

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

Explanation:

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

QUESTION 7

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 Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

QUESTION 8

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table1.

You need to identify the extent of the data skew in Table1.

What should you do in Synapse Studio?

- A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
- B. Connect to the built-in pool and run DBCC CHECKALLOC.
- C. Connect to Pool1 and run DBCC CHECKALLOC.
- D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Correct Answer: D

Explanation:

Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet>

QUESTION 9

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 two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Explanation:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note:

REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

QUESTION 10

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

```
CREATE TABLE dbo.Employee
(
    EmployeeID          INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName           VARCHAR(100) NOT NULL,
    LastName            VARCHAR(100) NOT NULL,
    Title               VARCHAR(100) NULL,
    LastHireDate         DATETIME NULL,
    StreetAddress1      VARCHAR(500) NOT NULL,
    StreetAddress2      VARCHAR(500) NOT NULL,
    StreetAddress3      VARCHAR(500) NOT NULL,
    City                VARCHAR(200) NOT NULL,
    StateName            VARCHAR(20) NOT NULL,
    Salary              VARCHAR(20) NULL,
    PhoneNumber         VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance.

Which two actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Change Salary to the money data type.
- B. Change PhoneNumber to the float data type.
- C. Change LastHireDate to the datetime2(7) data type.
- D. Change PhoneNumber to the bigint data type.
- E. Change LastHireDate to the date data type.

Correct Answer: AE

QUESTION 11

HOTSPOT

You have a Microsoft SQL Server 2017 server that hosts five databases.

You Plan to migrate the databases to Azure.

You need to recommend a solution that meets the following requirements:

- Automatically scales compute based on the workload demand
- Provides per-second billing

What should you include in the recommendation? To answer, select the appropriate options in the answer area.

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

Azure service:

- A single Azure SQL database in the provisioned compute tier
- A single Azure SQL database in the serverless compute tier
- An Azure SQL Database elastic pool
- Azure SQL Managed Instance

Service tier:

- Basic
- General Purpose
- Standard

Correct Answer:

Azure service:

- A single Azure SQL database in the provisioned compute tier
- A single Azure SQL database in the serverless compute tier
- An Azure SQL Database elastic pool
- Azure SQL Managed Instance

Service tier:

- Basic
- General Purpose
- Standard

QUESTION 12

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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query `sys.dm_exec_requests` and discover that the wait type is `PAGELATCH_UP` and the wait_resource is `2:3:905856`.

You need to improve system performance.

Solution: You change the data file for the master database to autogrow by 10 percent.

Does this meet the goal?

[DP-300 Exam Dumps](#) [DP-300 PDF Dumps](#) [DP-300 VCE Dumps](#) [DP-300 Q&As](#)

<https://www.ensurepass.com/DP-300.html>

[Download Full Version DP-300 Exam Dumps\(Updated in Feb/2023\)](#)

- A. Yes
- B. No

Correct Answer: B

Explanation:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

QUESTION 13

You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.

User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

What should you do?

- A. Create User1 in the master database.
- B. Assign User1 the db_datareader role for the master database.
- C. Assign User1 the db_datareader role for the databases that User1 creates.
- D. Grant select on sys.databases to public in the master database.

Correct Answer: A

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage>

QUESTION 14

HOTSPOT

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

You need to initiate a backup of Sales.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.