DP-300 Q&As

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/dp-300.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

HOTSPOT

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table.

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.

Hot Area:

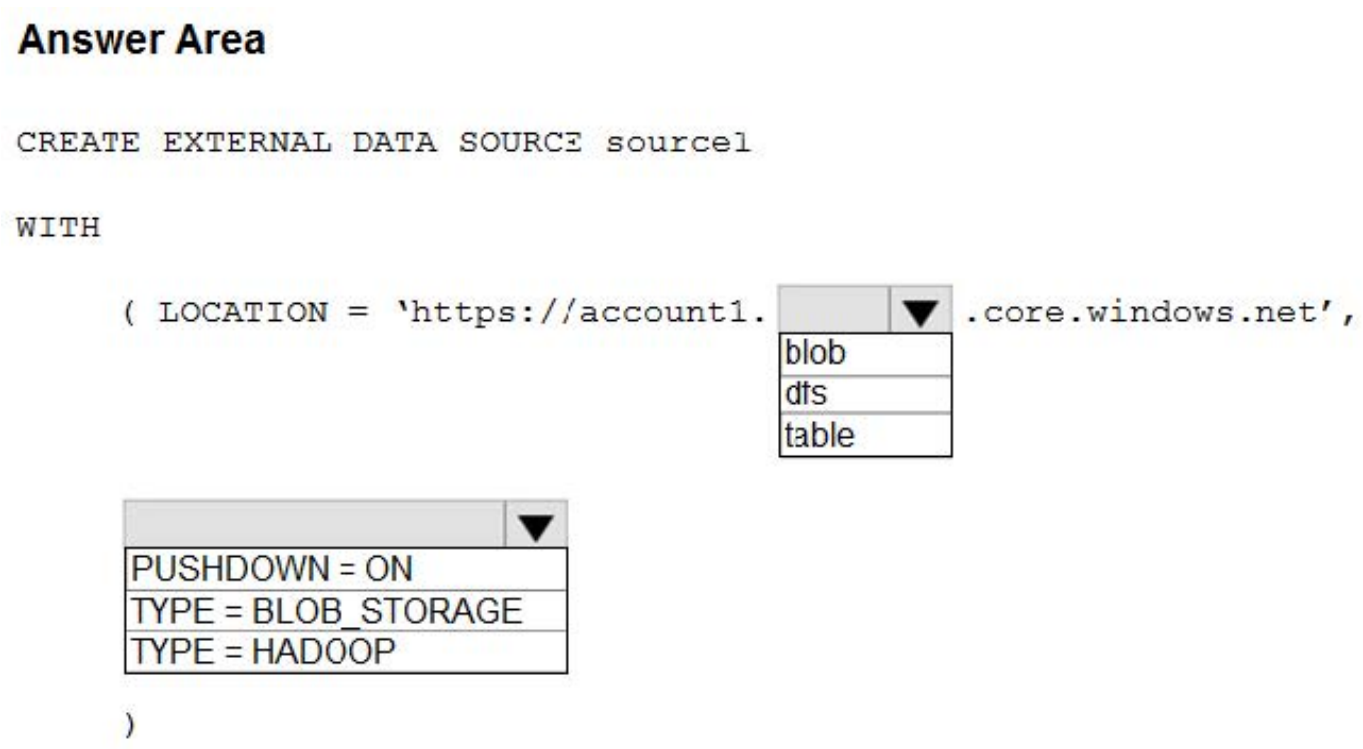

Correct Answer:

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1

WITH

      ( LOCATION = 'https://account1.  [ blob / dfs / table ]  .core.windows.net',

        [ PUSHDOWN = ON / TYPE = BLOB_STORAGE / TYPE = HADOOP ]

        )
```

Box 1: blob The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi WITH (LOCATION = \\'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/\\',

TYPE = HADOOP)

Box 2: HADOOP

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

QUESTION 2

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.

You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

A. Performance Recommendations in the Azure portal

B. Extended Events in Microsoft SQL Server Management Studio (SSMS)

C. Query Store in Microsoft SQL Server Management Studio (SSMS)

D. Query Performance Insight in the Azure portal

Correct Answer: C

Use the Query Store Page in SQL Server Management Studio.

Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by

using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using- the-query-store

QUESTION 3

You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features.

You plan to migrate to Azure SQL.

Which service should you use?

A. Azure SQL Database

B. SQL Server on an Azure Virtual Machine

C. Azure SQL Managed Instance

D. Azure Database for MySQL

Correct Answer: B

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-overview

QUESTION 4

You have an Azure subscription that contains 20 Azure SQL databases.

You create a Transact-SQL statement to perform index maintenance on a database.

You need to schedule the statement to run once daily against each database by using Transact-SQL commands.

What should you use to schedule the statement?

A. an Azure function

B. a SQL Server Agent Job

C. an elastic job

D. Azure Automation

Correct Answer: C

QUESTION 5

HOTSPOT

You are building a database in an Azure Synapse Analytics serverless SQL pool.

You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container.

Records are structured as shown in the following sample.

```
{
    "id":123,
    "address_housenumber": "19c",
    "address_line1": "Memory Lane",
    "applicant1_name": "Jane",
    "applicant2_name": "Dev"
}
```

The records contain two applicants at most.

You need to build a table that includes only the address fields.

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.

Hot Area:

Answer Area

```
[ CREATE EXTERNAL TABLE | CREATE TABLE | CREATE VIEW ] applications
WITH (
      LOCATION = 'applications/',
      DATA_SOURCE = applications_ds,
      FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
[ CROSS APPLY | OPENJSON | OPENROWSET ] (BULK 'https://contoso1.dfs.core.windows.net/
applications/year=*/*.parquet',
      FORMAT = 'PARQUET') AS [r]
GO
```

Correct Answer:

Answer Area

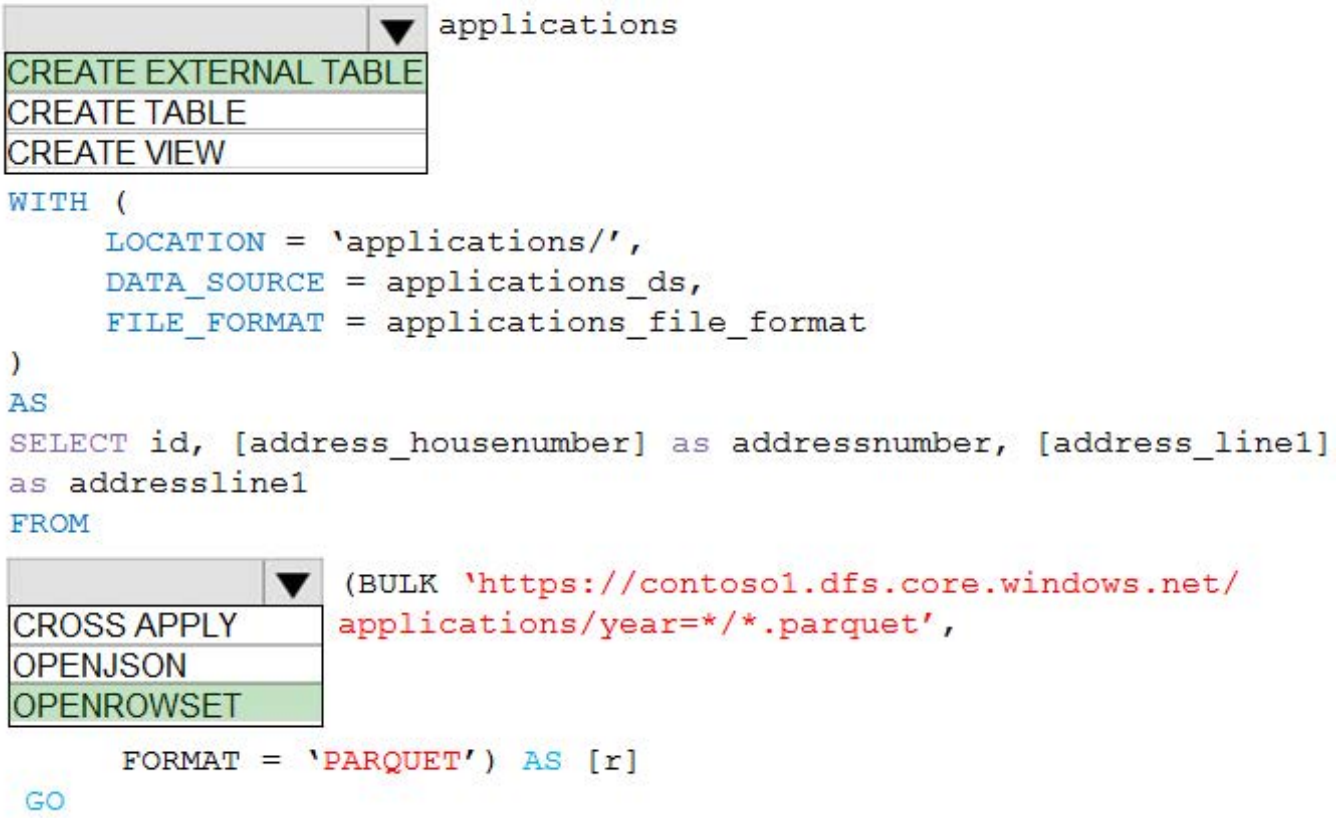

Box 1: CREATE EXTERNAL TABLE An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.

Syntax:

CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name } ([,...n]) WITH (

LOCATION = \\'folder_or_filepath\\',

DATA_SOURCE = external_data_source_name,

FILE_FORMAT = external_file_format_name

Box 2. OPENROWSET

When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example:

AS

SELECT decennialTime, stateName, SUM(population) AS population

FROM

OPENROWSET(BULK \\'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/*.parquet\\',

FORMAT=\\'PARQUET\\') AS [r] GROUP BY decennialTime, stateName GO

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables