



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**

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements. What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table the has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

Correct Answer: B

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-identity>

QUESTION 2

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL1 to Azure. The solution must meet the following requirements:

1.
The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

2.

Administrative effort must be minimized. What should you use to host the databases?

- A. a single Azure SQL database
- B. an Azure SQL Database elastic pool
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed Instance

Correct Answer: D

Linked servers enable the SQL Server Database Engine and Azure SQL Managed Instance to read data from the remote data sources and execute commands against the remote database servers (for example, OLE DB data sources) outside of the instance of SQL Server.



Linked servers are available in SQL Server Database Engine and Azure SQL Managed Instance. They are not enabled in Azure SQL Database singleton and elastic pools.

Reference: <https://learn.microsoft.com/en-us/sql/relational-databases/linked-servers/linked-servers-database-engine>

QUESTION 3

What should you use to migrate the PostgreSQL database?

- A. Azure Data Box
- B. AzCopy
- C. Azure Database Migration Service
- D. Azure Site Recovery

Correct Answer: C

Reference: <https://docs.microsoft.com/en-us/azure/dms/dms-overview>

QUESTION 4

You have an Azure SQL database named DB1.

A user named User 1 has an Azure AD account.

You need to provide User1 with the ability to add and remove columns from the tables in DB1. The solution must use the principle of least privilege.

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

NOTE: Each correct selection is worth one point

- A. Assign the database user the db_owner role.
- B. Create a contained database user.
- C. Create a login and an associated database user.
- D. Assign the database user the db_ddladmin role.

Correct Answer: BD

- B. Create a contained database user.
 - D. Assign the database user the db_ddladmin role.
-

QUESTION 5

You have an Azure Synapse Analytics Apache Spark pool named Pool1.



You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.

You need to load the files into the tables. The solution must maintain the source data types.

What should you do?

- A. Load the data by using PySpark.
- B. Load the data by using the OPENROWSET Transact-SQL command in an Azure Synapse Analytics serverless SQL pool.
- C. Use a Get Metadata activity in Azure Data Factory.
- D. Use a Conditional Split transformation in an Azure Synapse data flow.

Correct Answer: B

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools. Serverless SQL pool enables you to query data

in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage>

[Latest DP-300 Dumps](#)

[DP-300 VCE Dumps](#)

[DP-300 Exam Questions](#)