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.pass2lead.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

2 / 6

QUESTION 1

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database.

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

NOTE: Each correct selection is worth one point.

A. Implement row-level security (RLS).

B. Create users in each database.

C. Configure the database firewall.

D. Configure the server firewall.

E. Create logins in the master database.

F. Implement Always Encrypted.

Correct Answer: BE

E: Generating Logins

Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:

CREATE LOGIN readonlylogin WITH password=\\'1231!#ASDF!a\\';

You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.

B: Creating Users.

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a

user:

CREATE USER readonlyuser FROM LOGIN readonlylogin;

Incorrect Answers:

Not C: Database-level firewall rules only apply to individual databases. However, we need to create logins and users.

Not D: Server-level IP firewall rules apply to all databases within the same server.

Reference: https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/

QUESTION 2

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 to 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

QUESTION 3

HOTSPOT

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.
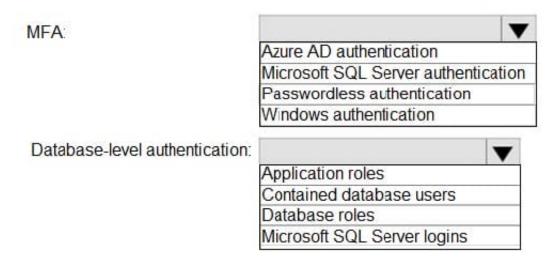
Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

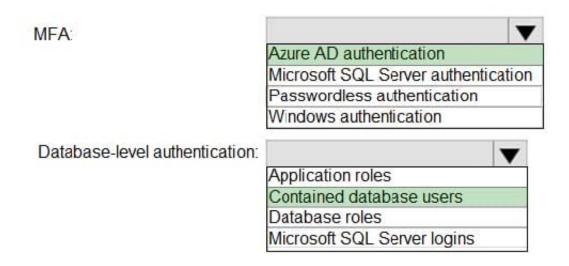NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

MFA:

Azure AD authentication
Microsoft SQL Server authentication
Passwordless authentication
Windows authentication

Database-level authentication:

Application roles
Contained database users
Database roles
Microsoft SQL Server logins

Correct Answer:

Answer Area

MFA:

Azure AD authentication
Microsoft SQL Server authentication
Passwordless authentication
Windows authentication

Database-level authentication:

Application roles
Contained database users
Database roles
Microsoft SQL Server logins

Box 1: Azure AD authentication

Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.

Box 2: Contained database users

Azure Active Directory Uses contained database users to authenticate identities at the database level.

Incorrect:

SQL authentication: To connect to dedicated SQL pool (formerly SQL DW), you must provide the following information:

1.

Fully qualified servername

2.

Specify SQL authentication

3.

Username

4.

Password

5.

Default database (optional)

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-authentication

---

QUESTION 4

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

DB1 contains a fact table named Table.

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

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 5

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1.

You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

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

NOTE: Each correct selection is worth one point.

A. From the properties of DB1, configure geo-replication.

B. From the properties of Server1, add a failover group.

C. Create a new Azure SQL Database server named Server2.

D. From the properties of Server1, configure retention for DB1.

E. Create a new Azure SQL Database instance named DB2.

Correct Answer: BC

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell#best-practices-for-sql-database https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azure-portal

[DP-300 Practice Test](#)   [DP-300 Exam Questions](#)   [DP-300 Braindumps](#)