

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



QUESTION 1

You have an Azure subscription.

You need to deploy a new Azure SQL database by using Azure Command-Line Interface (CLI).

Which three parameters are required?

- A. --name, --edition, and --capacity
- B. --name, --tier, and --min-capacity
- C. --name, --resource-group, and --server
- D. --name, --licence-type, and --capacity

Correct Answer: C

az sql db create Create a database

Syntax: az sql db create --name --resource-group --server

Plus many additional optional parameters.

Required Parameters

--name -n

Name of the Azure SQL Database.

--resource-group -g

Name of resource group. You can configure the default group using az configure --defaults group=.

--server -s

Name of the Azure SQL Server. You can configure the default using az configure --defaults sql-server=.

Example:

```
echo "Creating $database on $server..."
```

```
az sql db create --resource-group $resourceGroup --server $server --name $database --sample-name AdventureWorksLT --edition GeneralPurpose --family Gen5 --capacity 2 --zone-redundant true # zone redundancy is only supported on
```

premium and business critical service tiers

Reference:

<https://learn.microsoft.com/en-us/cli/azure/sql/db>

QUESTION 2

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 Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead use a serverless SQL pool to create an external table with the extra column.

Note: In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables>

QUESTION 3

You are modifying an existing disaster recovery solution for an Azure SQL managed instance that contains a failover group named FG1.

You need to ensure the maximum in-transit time for FG1 when an automatic failover occurs.

What should you configure?

A. an availability group

B. a secondary managed instance

C. a failover policy

D. a grace period

Correct Answer: D

Configuring Auto-Failover group on Azure SQL server Select the read/write failover policy. Select Automatic if you need automatic failover when the primary goes down. Select the read-write grace period. As the data is replicated asynchronously, the immediate failover may result in data loss. So, we must set the grace period accordingly to reduce data loss. The grace period is the time the SQL service waits before triggering automatic failover

when an outage occurs.

Reference:

<https://www.sqlshack.com/azure-sql-server-auto-failover-groups>

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: CD

C: Logins and users: A user account in a database can be associated with a login that is stored in the master database or can be a user name that is stored in an individual database.

A login is an individual account in the master database, to which a user account in one or more databases can be linked. With a login, the credential information for the user account is stored with the login.

A user account is an individual account in any database that may be, but does not have to be, linked to a login. With a user account that is not linked to a login, the credential information is stored with the user account.

D: db_ddladmin: Members of the db_ddladmin fixed database role can run any Data Definition Language (DDL) command (such as adding and removing columns) in a database.

Reference:

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

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles?view=sql-server-ver15>

QUESTION 5

You have an Azure subscription.

You create a logical SQL server that hosts four databases. Each database will be used by a separate customer.

You need to ensure that each customer can access only its own database. The solution must minimize administrative effort.

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

NOTE: Each correct selection is worth one point.

- A. Deny public access.
- B. Create a private endpoint.
- C. Create a database-level firewall rule.
- D. Create a network security group (NSG).
- E. Create a server-level firewall rule.

Correct Answer: B

You can connect to an Azure SQL server using an Azure Private Endpoint.

Azure Private endpoint is the fundamental building block for Private Link in Azure. It enables Azure resources, like virtual machines (VMs), to privately and securely communicate with Private Link resources such as Azure SQL server.

Reference:

<https://learn.microsoft.com/en-us/azure/private-link/tutorial-private-endpoint-sql-portal>

[Latest DP-300 Dumps](#)

[DP-300 Practice Test](#)

[DP-300 Braindumps](#)