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.geekcert.com/dp-300.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1.

What should you do?

A. Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property.

B. Use Transact-SQL to enable the force_last_good_plan option.

C. From the Azure portal, configure automatic tuning.

Correct Answer: B

For Azure SQL Managed Instance, the supported option FORCE_LAST_GOOD_PLAN can only be configured through T-SQL.

Enable automatic tuning on server On the server level you can choose to inherit automatic tuning configuration from "Azure Defaults" or not to inherit the configuration. Azure defaults are FORCE_LAST_GOOD_PLAN enabled, CREATE_INDEX disabled, and DROP_INDEX disabled.

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable

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 SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

1.

During normal operations, provide at least two readable copies of Sales.

2.

Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

QUESTION 3

HOTSPOT

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To add read-only replicas of DB1:

- Create a replica on the same logical server.
- Create a new logical server and configure geo-replication.
- Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

- Add an ApplicationIntent entry to the connection string.
- Add a MultiSubnetFailover entry to the App1 connection string.
- Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Correct Answer:

Answer Area

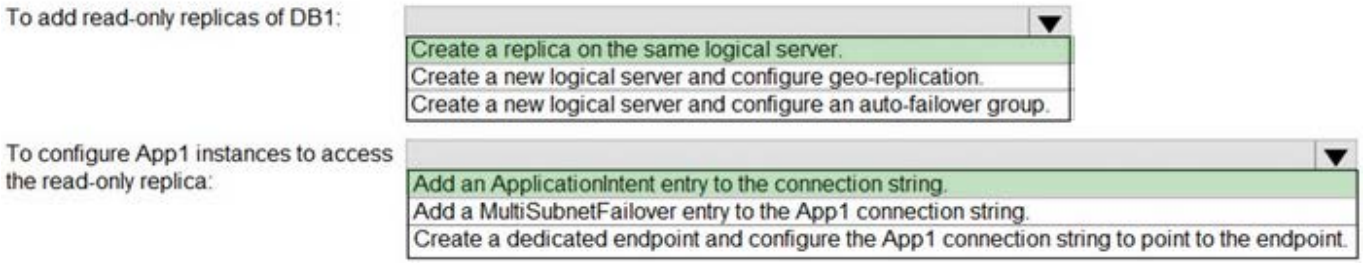

Reference: https://sqlserverguides.com/read-only-replica-azure-sql/

QUESTION 4

You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:

1.

Support real-time data replication to a different geographic region.

2.

Use Azure as a disaster recovery target.

3.

Minimize costs and administrative effort. What should you include in the recommendation?

A. database mirroring on an instance of SQL Server on Azure Virtual Machines

B. availability groups for SQL Server on Azure Virtual Machines

C. an Azure SQL Managed Instance link

D. transactional replication to an Azure SQL Managed Instance

Correct Answer: D

Transactional replication is a feature of Azure SQL Managed Instance and SQL Server that enables you to replicate data from a table in Azure SQL Managed Instance or a SQL Server instance to tables placed on remote databases. This feature allows you to synchronize multiple tables in different databases.D

You can use transactional replication to push changes made in an Azure SQL Managed Instance to:

1.

A SQL Server database - on-premises or on Azure VM

2.

A database in Azure SQL Database

3.

An instance database in Azure SQL Managed Instance

Incorrect:

Not A: Operates within the same location.

Database mirroring is a solution for increasing the availability of a SQL Server database.

Database mirroring maintains two copies of a single database that must reside on different server instances of SQL Server Database Engine. Typically, these server instances reside on computers in different locations.

Not B: Always On availability group on SQL Server on Azure VMs - only within a single region

Always On availability groups on Azure Virtual Machines are similar to Always On availability groups on-premises, and rely on the underlying Windows Server Failover Cluster.

VM redundancy

To increase redundancy and high availability, SQL Server VMs should either be in the same availability set, or different availability zones.

Placing a set of VMs in the same availability set protects from outages within a data center caused by equipment failure (VMs within an Availability Set do not share resources) or from updates (VMs within an availability set are not updated at

the same time).

Availability Zones protect against the failure of an entire data center, with each Zone representing a set of data centers within a region.

Reference:

https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/replication-transactional-overview

https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-overview

QUESTION 5

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 copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.

Reference: https://docs.microsoft.com/en-US/azure/data-factory/transform-data