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.passapply.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 manage 100 Azure SQL managed instances located across 10 Azure regions.

You need to receive voice message notifications when a maintenance event affects any of the 10 regions.

The solution must minimize administrative effort.

What should you do?

A. From the Azure portal, create a service health alert.

B. From the Azure portal, create an Azure Advisor operational excellence alert.

C. From Microsoft SQL Server Management Studio (SSMS), configure a SQL Server agent job.

D. From the Azure portal, configure an activity log alert.

Correct Answer: C

QUESTION 2

What should you implement to meet the disaster recovery requirements for the PaaS solution?

A. Availability Zones

B. failover groups

C. Always On availability groups

D. geo-replication

Correct Answer: B

Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active georeplication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.

The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance

availability in the primary region.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

QUESTION 3

You have a version-8.0 Azure Database for MySQL database.

You need to identify which database queries consume the most resources.

Which tool should you use?

A. Query Store

B. Metrics

C. Query Performance Insight

D. Alerts

Correct Answer: A

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns. Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance.

Reference: https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store

QUESTION 4

You have an Azure subscription that contains an Azure SQL database. The database contains a table named tablet that uses partitioned columnstores. You need to configure table1 to meet the following requirements:

1.

Each partition must be compressed.

2.

The compression ratio must be maximized.

3.

You must be able to index the compressed data. What should you use?

A. page compression

B. columnstore compression

C. GZIP compression

D. columnstore archival compression

Correct Answer: D

SQL Server, Azure SQL Database, and Azure SQL Managed Instance support row and page compression for rowstore tables and indexes, and support columnstore and columnstore archival compression for columnstore tables and

indexes.

For columnstore tables and indexes, all columnstore tables and indexes always use columnstore compression and this is not user configurable.

Compressing columnstore indexes with archival compression, causes the index to perform slower than columnstore indexes that do not have the archival compression. Use archival compression only when you can afford to use extra time and

CPU resources to compress and retrieve the data.

The benefit of archival compression, is reduced storage, which is useful for data that is not accessed frequently. For example, if you have a partition for each month of data, and most of your activity is for the most recent months, you could archive older months to reduce the storage requirements. Reference: https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression

QUESTION 5

You have 10 Azure virtual machines that have SQL Server installed.

You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups.

What should you include in the backup strategy?

A. Automated Backup in the SQL virtual machine settings

B. Azure Backup

C. Azure Site Recovery

D. SQL Server Agent jobs

Correct Answer: B

Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault. There are several advantages that this solution provides, especially for Enterprises.

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/backup-restore#azbackup