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.leads4pass.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 are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year. During peak usage, the database will require the following:

1.

24 cores

2.

500 GB of storage

3.

124 GB of memory

4.

More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard.

Which service tier should you use during peak usage?

A. Business Critical

B. Premium

C. Hyperscale

Correct Answer: A

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcore-single-databases#business-critical---provisioned-compute---gen4

QUESTION 2

You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.

You need to implement alerts by using Azure Monitor. The solution must meet the following requirements:

1.

Minimize costs.

2.

Aggregate Intelligent Insights telemetry from each database. What should you do?

A. From the Diagnostic settings of each database, select Send to Log Analytics.

B. From the Diagnostic settings of each database, select Stream to an event hub.

C. From the Diagnostic settings of SQLMI1, select Send to Log Analytics.

D. From the Diagnostic settings of SQLMI1, select Stream to an event hub.

Correct Answer: A

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry

QUESTION 3

You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

Name	Type
Column1	Ntext
Column2	Geometry
Column3	Image
Column4	Varchar
Column5	Datetime2

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point

A. Column1

B. Column2

C. Column3

D. Column4

E. Column5

Correct Answer: AD

Ntext and Varchar columns can be encrypted.

QUESTION 4

You have an instance of SQL Server on Azure Virtual Machine named SQL1.

You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.

Where should you store the metrics?

A. a Log Analytics workspace

B. Azure Event Hubs

C. Azure SQL Database

D. an Azure Blob storage container

Correct Answer: A

QUESTION 5

You deploy an instance of SQL Server on Azure Virtual Machines: named SQL1 that hosts multiple databases.

You configure the full recovery model for all the databases.

You perform a full backup of the master database on SQL1.

You need to a perform an additional backup of the master database on SQL1. The solution must minimize how long it takes to perform the backup.

Which type of backup should you perform?

A. log

B. full

C. differential

D. tail-log

Correct Answer: C

Under the full recovery model, using differential backups can reduce the number of log backups that you have to restore.

Benefits

Creating a differential backup can be much faster than creating a full backup. A differential backup records only the data that has changed since the full backup upon the differential backup is based.

Note:

A differential backup is based on the most recent, previous full data backup. A differential backup captures only the data that has changed since that full backup. The full backup upon which a differential backup is based is known as the base

of the differential.

Incorrect:

* tail-log A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the

database.

Reference: https://learn.microsoft.com/en-us/sql/relational-databases/backup-restore/differential-backups-sql-server