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.pass4itsure.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 Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.

Data to be loaded is identified by a column named LastUpdatedDate in the source table.

You plan to execute the pipeline every four hours.

You need to ensure that the pipeline execution meets the following requirements:

1.

Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits.

2.

Supports backfilling existing data in the table. Which type of trigger should you use?

A. tumbling window

B. on-demand

C. event

D. schedule

Correct Answer: A

The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past. Incorrect Answers:

D: Schedule trigger does not support backfill scenarios. Pipeline runs can be executed only on time periods from the current time and the future.

Reference: https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers

QUESTION 2

You have an Azure Data Factory pipeline that is triggered hourly.

The pipeline has had 100% success for the past seven days.

The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

```
ErrorCode=UserErrorFileNotFound,
'Type=Microsoft.DataTransfer.Common.Shared.HybridDeliveryException,Message=ADLS
Gen2 operation failed for: Operation returned an invalid status code
'NotFound'. Account: 'contosoproduksouth' FileSystem: wwi.Path:
'BIKES/CARBON/year=2021/month=01/day=10/hour=06'. ErrorCode:
'PathNotFound'.Message: 'The specified path does not exist.'. RequestId:
'6d269b78-901f-001b-4924-e7a7bc000000'. TimeStamp: 'Sun, 10 Jan 2021 07:45:05
```

What is a possible cause of the error?

A. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.

B. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.

C. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.

D. The pipeline was triggered too early.

Correct Answer: A

QUESTION 3

You have an Azure subscription that contains three instances of SQL Server on Azure Virtual Machines.

You plan to implement a disaster recovery solution.

You need to be able to perform disaster recovery drills regularly. The solution must meet the following requirements:

1.

Minimize administrative effort for the recovery drills.

2.

Isolate the recovery environment from the production environment What should you use?

A. native Microsoft SQL Server backup

B. Azure Site Recovery

C. Recovery Services vaults

D. Azure Backup

Correct Answer: B

Set up disaster recovery for SQL Server

You can protect the SQL Server back end of an application. You do so by using a combination of SQL Server business continuity and disaster recovery (BCDR) technologies and Azure Site Recovery.

SQL Server disaster recovery capabilities include:

Failover clustering

Always On availability groups

Database mirroring

Log shipping

Active geo-replication

Auto-failover groups

Note: Azure Recovery Services contributes to your BCDR strategy:

Site Recovery service: Site Recovery helps ensure business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary site

to a secondary location. When an outage occurs at your primary site, you fail over to a secondary location, and access apps from there. After the primary location is running again, you can fail back to it.

Backup service: The Azure Backup service keeps your data safe and recoverable.

Site Recovery can manage replication for:

Azure VMs replicating between Azure regions Replication from Azure Public Multi-Access Edge Compute (MEC) to the region Replication between two Azure Public MECs On-premises VMs, Azure Stack VMs, and physical servers

Reference: https://learn.microsoft.com/en-us/azure/site-recovery/site-recovery-sql https://learn.microsoft.com/en-us/azure/site-recovery/site-recovery-overview

QUESTION 4

You have SQL Server on Azure virtual machines in an availability group.

You have a database named DB1 that is NOT in the availability group.

You create a full database backup of DB1.

You need to add DB1 to the availability group.

Which restore option should you use on the secondary replica?

A. Restore with Recovery

B. Restore with Norecovery

C. Restore with Standby

Correct Answer: B

Prepare a secondary database for an Always On availability group requires two steps:

1.

Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY

2.

Join the restored database to the availability group.

Reference:

https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a- secondary-database-for-an-availability-group-sql-server

QUESTION 5

HOTSPOT

You have an Azure SQL database.

You are reviewing a slow performing query as shown in the following exhibit.

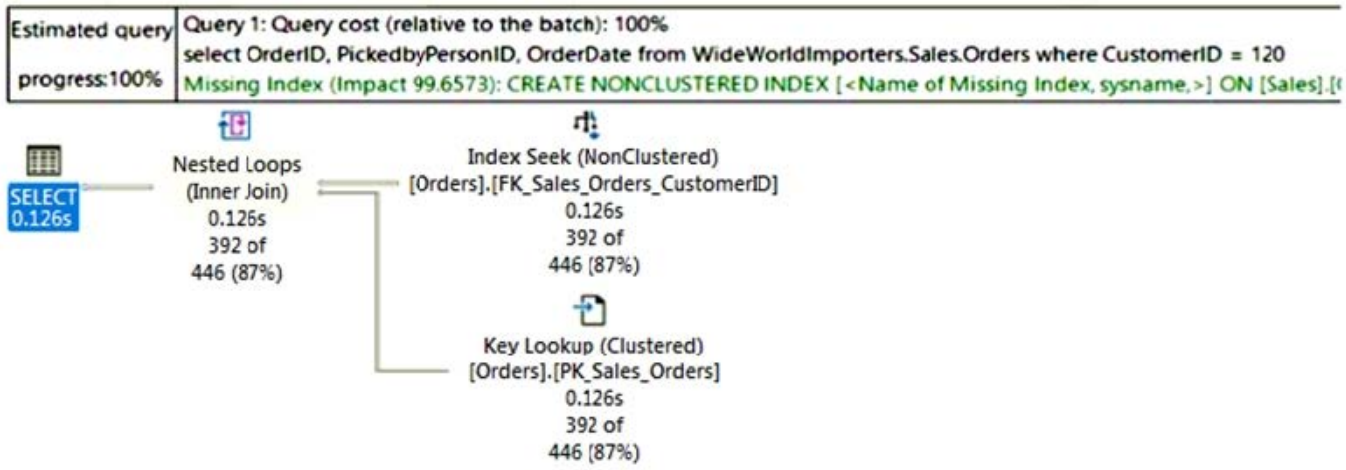

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

The exhibit shows **[answer choice]**.

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

Correct Answer:

Answer Area

The exhibit shows **[answer choice]**.

- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.

- Index Seek
- Key Lookup
- Nested Loops

Reference: https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15

QUESTION 6

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 7

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 two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

There is no REPLACE option in Azure SQL Database, or even there is no restore option with T-SQL. https://learn.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups?view=azuresql-dbandtabs=azure-portal

QUESTION 8

You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

1.

Automate the patching of SQL Server and Windows Server.

2.

Automate full database backups and transaction log backups of the databases on VM1.

3.

Minimize administrative effort. What should you do first?

A. Enable a system-assigned managed identity for VM1

B. Register VM1 to the Microsoft.Sql resource provider

C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1

D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Correct Answer: B

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaasExtension) runs on Azure virtual machines to automate administration tasks. The SQL Server

IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent- extension-automate-management

QUESTION 9

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate    DATETIME NULL,
    StreetAddress1  VARCHAR(500) NOT NULL,
    StreetAddress2  VARCHAR(500) NOT NULL,
    StreetAddress3  VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Change Salary to the money data type.

B. Change PhoneNumber to the float data type.

C. Change LastHireDate to the datetime2(7) data type.

D. Change PhoneNumber to the bigint data type.

E. Change LastHireDate to the date data type.

Correct Answer: AE

QUESTION 10

HOTSPOT

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.

You need to ensure that app1 can access db1. The solution must use the principle of least privilege.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Correct Answer:

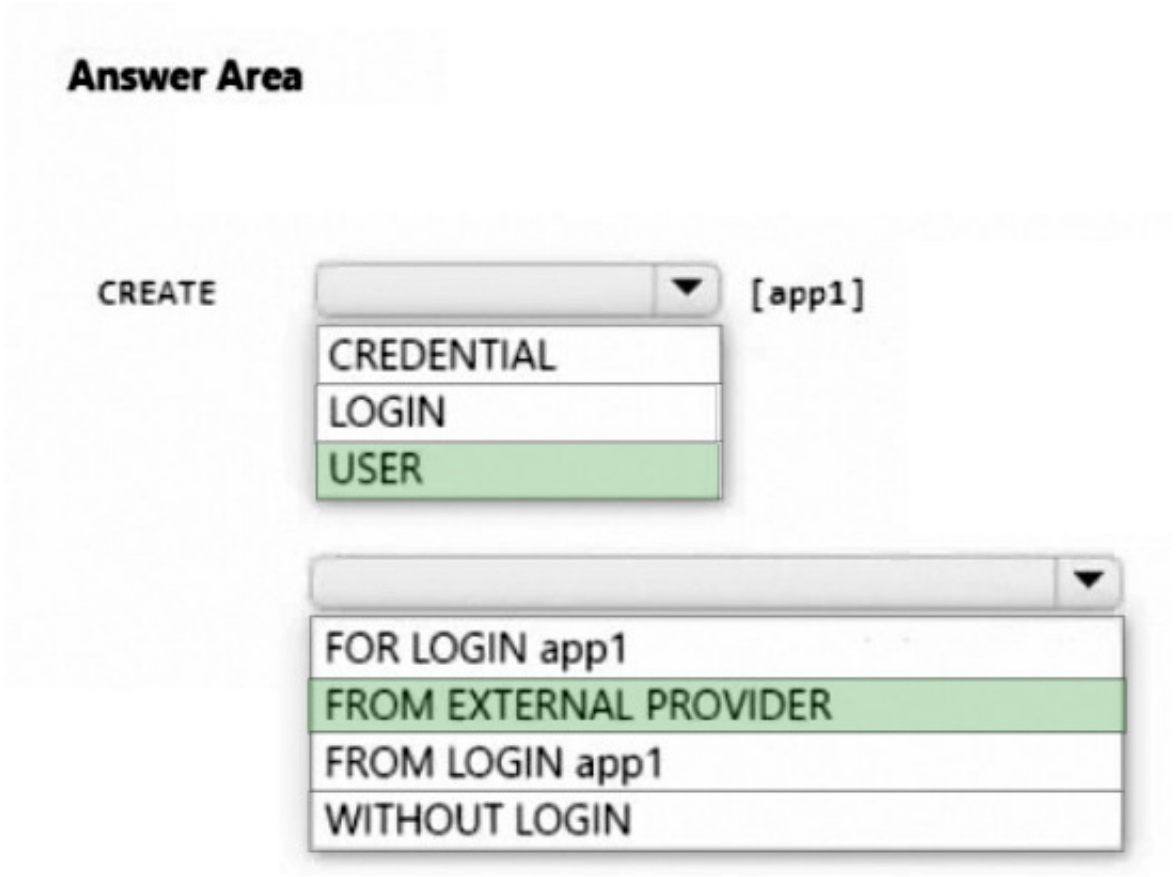

Box 1: USER

Azure Active Directory service principal with Azure SQL

When an Azure AD application is registered using the Azure portal or a PowerShell command, two objects are created in the Azure AD tenant:

An application object

A service principal object

SQL Database and SQL Managed Instance support the following Azure AD objects:

Azure AD users (managed, federated, and guest)

Azure AD groups (managed and federated)

Azure AD applications

The T-SQL command CREATE USER [Azure_AD_Object] FROM EXTERNAL PROVIDER on behalf of an Azure AD application is now supported for SQL Database.

Box 2: FROM EXTERNAL PROVIDER

Reference:

https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal

QUESTION 11

DRAG DROP

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change.

You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

Which three backups should you restore in sequence? To answer, move the appropriate backups from the list of backups to the answer area and arrange them in the correct order.

Select and Place:

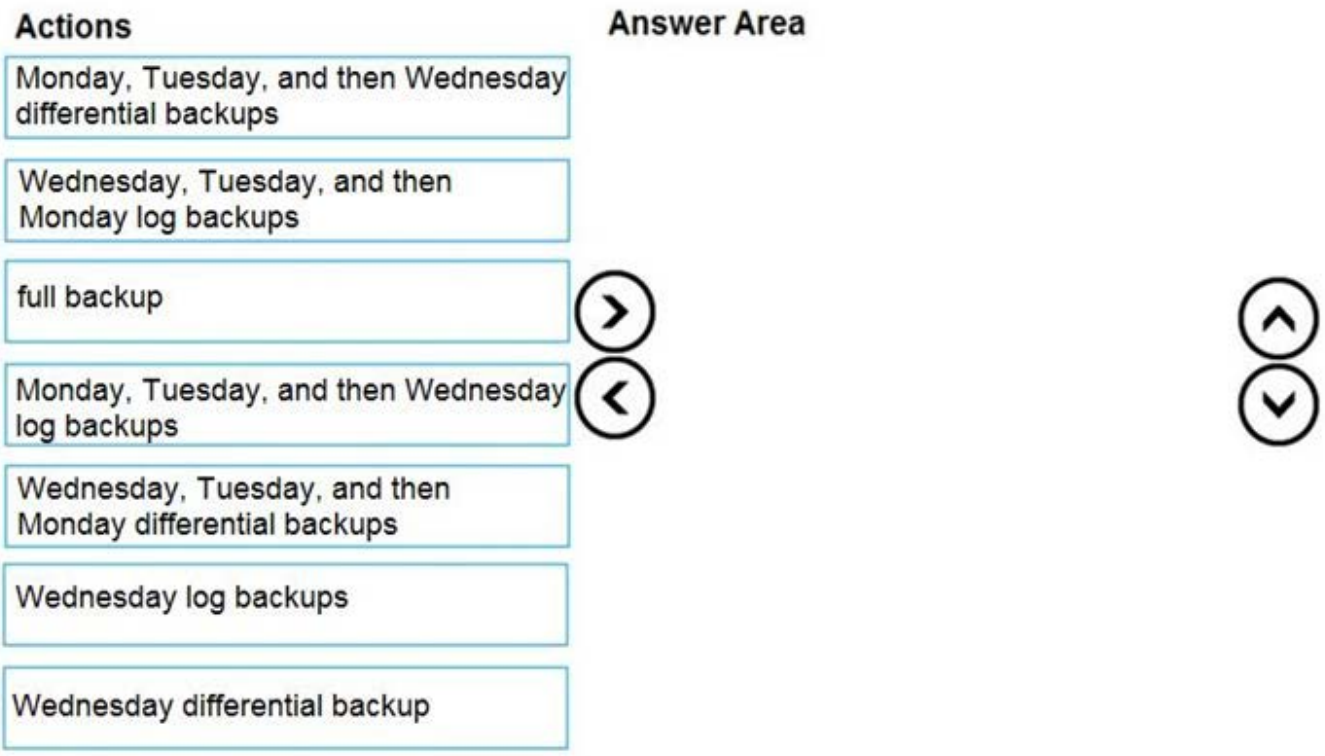

Correct Answer:

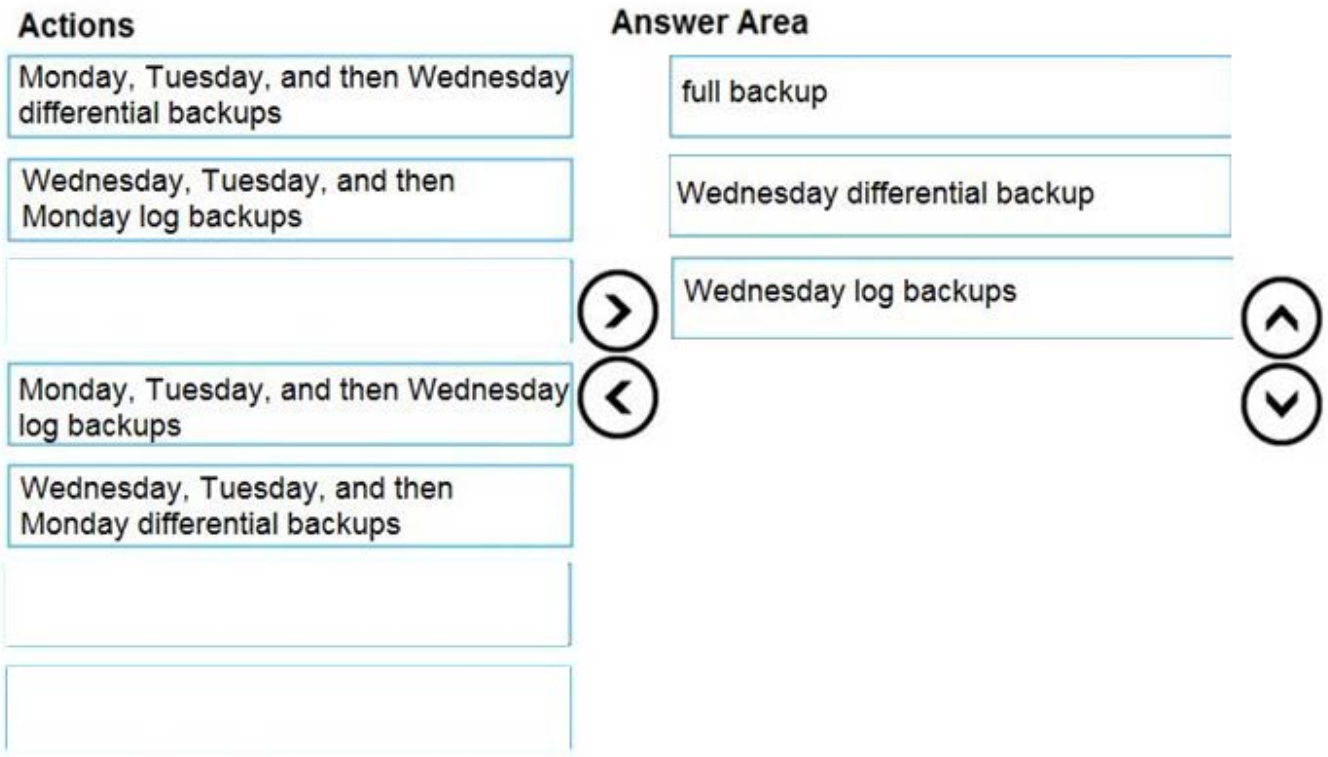

QUESTION 12

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8.

You need to deploy a failover cluster instance (FCI) to VM1 and VM2 that will use Azure shared disks. The solution must maximize resiliency.

Which quorum option should you use?

A. node majority with a cloud witness

B. node majority with no witness

C. node majority with a file share witness

D. node majority with a disk witness

Correct Answer: D

Configure quorum

Since the disk witness is the most resilient quorum option, and the FCI solution uses Azure shared disks, it\\'s recommended to configure a disk witness as the quorum solution.

Reference: https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/failover-cluster-instance-azure-shared-disks-manually-configure

QUESTION 13

You have an Azure subscription.

You need to deploy an Azure SQL database. The solution must meet the following requirements:

1.

Dynamically scale CPU resources.

2.

Ensure that the database can be paused to reduce costs.

What should you use?

A. the Business Critical service tier

B. the serverless compute tier

C. an elastic pool

D. the General Purpose service tier

Correct Answer: B

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns. The serverless compute tier is available in the General Purpose service tier and currently in preview in the Hyperscale service tier.

Reference: https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview

QUESTION 14

HOTSPOT

You have an Azure SQL database named DB1 that contains a table named Table1.

You run a query to load data into Table1.

The performance metrics of Table1 during the load operation are shown in the following exhibit.

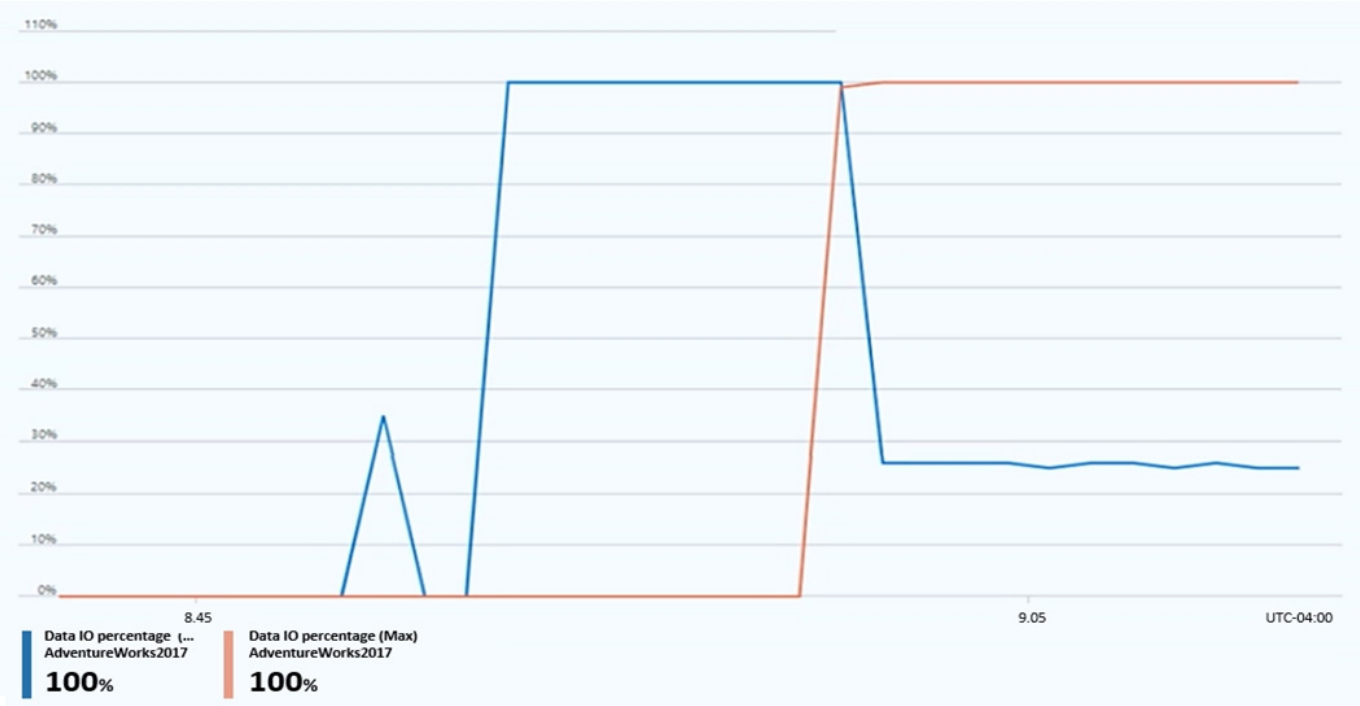

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To reduce how long it takes to complete the query you must **[answer choice]**.

scale the resource
use an elastic resource
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use **[answer choice]** table.

a temporary
an In-Memory OTLP durable
an In-Memory OTLP non durable

Correct Answer:

Answer Area

To reduce how long it takes to complete the query you must **[answer choice]**.

scale the resource
use an elastic resource
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use **[answer choice]** table.

a temporary
an In-Memory OTLP durable
an In-Memory OTLP non durable

Box 1: perform query tuning

Implementation considerations

Use memory-optimized tables for your core transaction tables, that is, the tables with the most performance-critical transactions. Use natively compiled stored procedures to optimize execution of the logic associated with the business

transaction. The more of the logic you can push down into stored procedures in the database, the more benefit you see from In-Memory OLTP.

To get started in an existing application:

Use the transaction performance analysis report to identify the objects you want to migrate.

Use the memory-optimization and native compilation advisors to help with migration.

Box 2: an In-Memory OLTP non-durable

Example: tempdb object replacement

Use non-durable tables and memory-optimized table types to replace your traditional tempdb based structures, such as temporary tables, table variables, and table-valued parameters (TVPs).

Memory-optimized table variables and non-durable tables typically reduce CPU and completely remove log IO, when compared with traditional table variables and #temp table.

Note: In-memory technologies enable you to improve performance of your application, and potentially reduce cost of your database.

In-Memory OLTP

In-Memory OLTP technology provides extremely fast data access operations by keeping all data in memory. It also uses specialized indexes, native compilation of queries, and latch-free data-access to improve performance of the OLTP

workload. There are two ways to organize your In-Memory OLTP data:

*

Memory-optimized rowstore format where every row is a separate memory object. This is a classic In-Memory OLTP format optimized for high-performance OLTP workloads. There are two types of memory-optimized tables that can be used in the memory-optimized rowstore format:

Durable tables (SCHEMA_AND_DATA) where the rows placed in memory are preserved after server restart. This type of tables behaves like a traditional rowstore table with the additional benefits of in-memory optimizations.

Non-durable tables (SCHEMA_ONLY) where the rows are not-preserved after restart. This type of table is designed for temporary data (for example, replacement of temp tables), or tables where you need to quickly load data before you move

it to some persisted table (so called staging tables).

*

Memory-optimized columnstore format where data is organized in a columnar format

Reference: https://learn.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/overview-and-usage-scenarios https://learn.microsoft.com/en-us/azure/azure-sql/in-memory-oltp-overview

QUESTION 15

You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.

User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.

You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

What should you do?

A. Create User1 in the master database.

B. Assign User1 the db_datareader role for the master database.

C. Assign User1 the db_datareader role for the databases that User1 creates.

D. Grant SELECTon sys.databases to public in the master database.

Correct Answer: A

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