

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

<https://www.2passeasy.com/dumps/DP-300/>



NEW QUESTION 1

- (Exam Topic 5)

You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1. You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context. 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.

```
EXECUTE AS 

|        |   |
|--------|---|
|        | ▼ |
| CALLER |   |
| LOGIN  |   |
| OWNER  |   |
| USER   |   |

 = 'user1@contoso.com'
```

GO

```
SELECT SUSER_SNAME ()
```

| | |
|----------|---|
| | ▼ |
| REVERT | |
| REVOKE | |
| ROLLBACK | |

GO

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15> <https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15>

NEW QUESTION 2

- (Exam Topic 5)

You have an Azure subscription that uses a domain named contoso.com. You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region. DBServer1 has a high volume of data changes and low latency requirements for data writes. You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2. What should you do?

- A. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous.
- B. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous.
- C. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous.
- D. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous.

Answer: C

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on>

NEW QUESTION 3

- (Exam Topic 5)

You have an Azure subscription. You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance. You need to configure the disks on the virtual machine. The solution must meet the following requirements:

- Minimize latency for transaction logs.
- Minimize the impact on IO Of the virtual machine.

Which type of disk should you use for each workload? To answer, drag the appropriate disk types to the correct workloads. Each disk type may be used once, at

all. You may need to drag the split bar between panes or scroll to view content.
 NOTE: Each correct selection is worth one point.

Disk Types

- Local
- Premium SSD
- Standard HDD
- Standard SSD
- Ultra Disk

Answer Area

TempDB:

Transaction logs:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 4

- (Exam Topic 5)

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux.

You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

To prepare for the migration:

To perform the migration, use:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To prepare for the migration:

To perform the migration, use:

NEW QUESTION 5

- (Exam Topic 5)

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.

You need to ensure that SOL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs.

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

Actions

- Enable pager notifications upon failure.
- Create a profile named application_dbmail_profile.
- Create a Database Mail account.
- Create a profile named AzureManagedInstance_dbmail_profile.
- Enable email notifications upon failure.

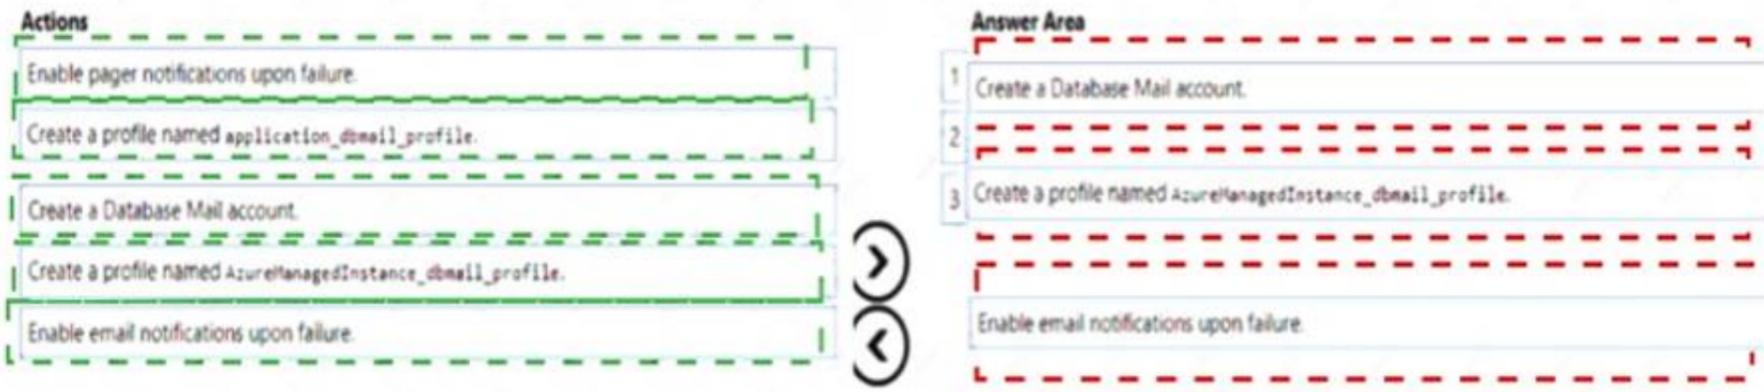
Answer Area

- 1
- 2
- 3

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 6

- (Exam Topic 5)

You have an Azure subscription.

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

- Dynamically scale CPU resources.
- 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

Answer: B

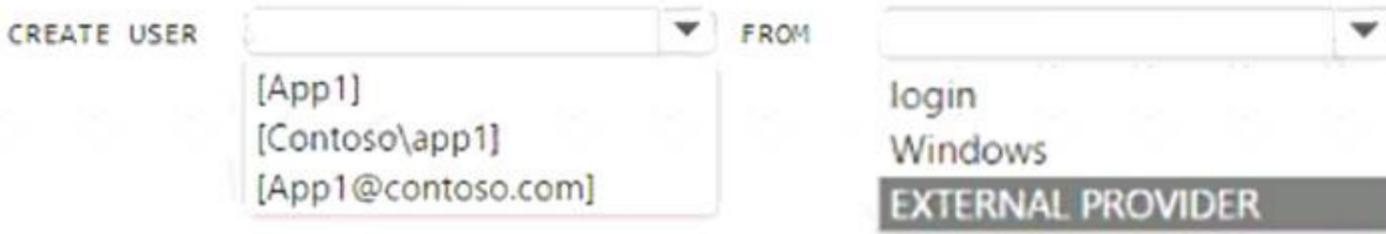
NEW QUESTION 7

- (Exam Topic 5)

You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.

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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<https://learn.microsoft.com/en-us/azure/app-service/tutorial-connect-msi-sql-database?tabs=windowsclient%2Ce>

NEW QUESTION 8

- (Exam Topic 5)

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2. You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

Answer: BD

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server>

NEW QUESTION 9

- (Exam Topic 5)

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:

- Ensure that the latest version of SQL Server is used.
- Support the SQL Server Agent service. Minimize administrative effort.

What should you use?

- A. SQL Server on Azure Virtual Machines
- B. Azure SQL Database
- C. an Azure SQL Database elastic pool
- D. Azure SQL Managed Instance

Answer: D

NEW QUESTION 10

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

- Capture raw event data and store the data in Azure Storage.
- Minimize the performance impact of capturing extended events.

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.

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
    ACTION (sqlserver.sql_text)
    WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
    event_file
    event_stream
    ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
    (MAX_MEMORY = 10 MB,
    EVENT_RETENTION_MODE=
    MAX_DISPATCH_LATENCY = 3 SEC
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
ACTION (sqlserver.sql_text)
WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0.
event_file
event_stream
ring_buffer
)
SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
(MAX_MEMORY = 10 MB,
EVENT_RETENTION_MODE=
ALLOW_MULTIPLE_EVENT_LOSS
ALLOW_SINGLE_EVENT_LOSS
NO_EVENT_LOSS
MAX_DISPATCH_LATENCY = 3 SECONDS)
```

NEW QUESTION 10

- (Exam Topic 5)

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

- > Send the output to an Azure Synapse.
- > Identify spikes and dips in time series data.
- > Minimize development and configuration effort.

Which should you include in the solution?

- A. Azure SQL Database
- B. Azure Databricks
- C. Azure Stream Analytics

Answer: C

Explanation:

Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics Reference:
<https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/> <https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output>

NEW QUESTION 14

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can perform the following tasks on DB1:

- Create jobs.
- View all jobs.
- Modify, delete, and disable the jobs the user created. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Location:
 DB1
 master
 msdb

Built-in role:
 SQLAgentOperatorRole
 SQLAgentReaderRole
 SQLAgentUserRole

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<https://learn.microsoft.com/en-us/sql/ssms/agent/sql-server-agent-fixed-database-roles?view=sql-server-ver16#s>

NEW QUESTION 18

- (Exam Topic 5)

You have an Azure SQL database named D61.

You need to identify how much unused space in megabytes was allocated to DB1.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
SELECT SUM(size/ 128.0) - CAST(FILEPROPERTY(name, 'SpaceUsed') AS Int)/128.0 AS DatabaseDataSpaceAllocatedUnusedInMB
FROM sys.databases
GROUP BY sys.databases
HAVING t: sys.dm_db_resource_stats
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
SELECT SUM(size/ 128.0) - CAST(FILEPROPERTY(name, 'SpaceUsed') AS Int)/128.0 AS DatabaseDataSpaceAllocatedUnusedInMB
FROM sys.databases
GROUP BY sys.databases
HAVING t: sys.dm_db_resource_stats
```

NEW QUESTION 19

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes. You need to ensure that indexes are created for the tables.

What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the dbcc dbreindex command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Answer: C

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview>

NEW QUESTION 22

- (Exam Topic 5)

You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. The solution must minimize downtime. What should you do?

- A. Configure Transaction Log Shipping.
- B. Implement Always On availability groups.
- C. Configure transactional replication.
- D. Import a BACPAC.

Answer: C

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1-migra>

NEW QUESTION 23

- (Exam Topic 5)

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases. The databases perform slower than expected. You need to identify whether the performance issue relates to the use of tempdb on sqlsrv1. What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Answer: D

Explanation:

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting. You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb> <https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-use-diagnostics-log>

NEW QUESTION 28

- (Exam Topic 5)

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

- > New data is accessed frequently and must be available as quickly as possible.
- > Data that is older than five years is accessed infrequently but must be available within one second when requested.
- > Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
- > Costs must be minimized while maintaining the required availability.

How should you manage the data? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Five-year-old data:

| | |
|--------------------------|---|
| | ▼ |
| Delete the blob. | |
| Move to archive storage. | |
| Move to cool storage. | |
| Move to hot storage. | |

Seven-year-old data:

| | |
|--------------------------|---|
| | ▼ |
| Delete the blob. | |
| Move to archive storage. | |
| Move to cool storage. | |
| Move to hot storage. | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery
 Older data not used frequently but expected to be available immediately when accessed
 Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot - Optimized for storing data that is accessed frequently.
 Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.
 Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.
 Box 2: Move to archive storage
 Example usage scenarios for the archive access tier include: Long-term backup, secondary backup, and archival datasets
 Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:
<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

NEW QUESTION 31

- (Exam Topic 5)

You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database. You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance. Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. General Purpose Standard
- B. Business Critical Premium
- C. Business Critical Memory Optimized Premium
- D. General Purpose Premium
- E. Business Critical Standard

Answer: BCD

NEW QUESTION 32

- (Exam Topic 5)

You are monitoring an Azure Stream Analytics job. You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero. You need to ensure that the job can handle all the events. What should you do?

- A. Remove any named consumer groups from the connection and use \$default.
- B. Change the compatibility level of the Stream Analytics job.
- C. Create an additional output stream for the existing input stream.
- D. Increase the number of streaming units (SUs).

Answer: D

Explanation:

Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring>

NEW QUESTION 37

- (Exam Topic 5)

You plan to migrate on-premises Microsoft SQL Server databases to Azure. You need to identify which deployment and resiliency options meet the following requirements:

- > Support user-initiated backups.
- > Support multiple automatically replicated instances across Azure regions.
- > Minimize administrative effort to implement and maintain business continuity. What should you identify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Deployment option:

| | |
|---------------------------------------|---|
| | ▼ |
| Azure SQL Managed Instance | |
| SQL Server on Azure Virtual Machines | |
| An Azure SQL Database single database | |

Resiliency option:

| | |
|---------------------------|---|
| | ▼ |
| Auto-failover group | |
| Active geo-replication | |
| Zone-redundant deployment | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SQL Server on Azure VMs

SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.

Box 2: Active geo-replication

Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-i> <https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure>

NEW QUESTION 38

- (Exam Topic 5)

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource. Which Databricks services should you log?

- A. clusters
- B. jobs
- C. DBFS
- D. SSH
- E. workspace

Answer: E

Explanation:

Cloud Provider Infrastructure Logs.

Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace. Customers, especially in the regulated industries, also need records on activities like:

- > User access control to cloud data storage
- > Cloud Identity and Access Management roles
- > User access to cloud network and compute

Azure Databricks offers three distinct workloads on several VM Instances tailored for your data analytics workflow—the Jobs Compute and Jobs Light Compute workloads make it easy for data engineers to build and execute jobs, and the All-Purpose Compute workload makes it easy for data scientists to explore, visualize, manipulate, and share data and insights interactively.

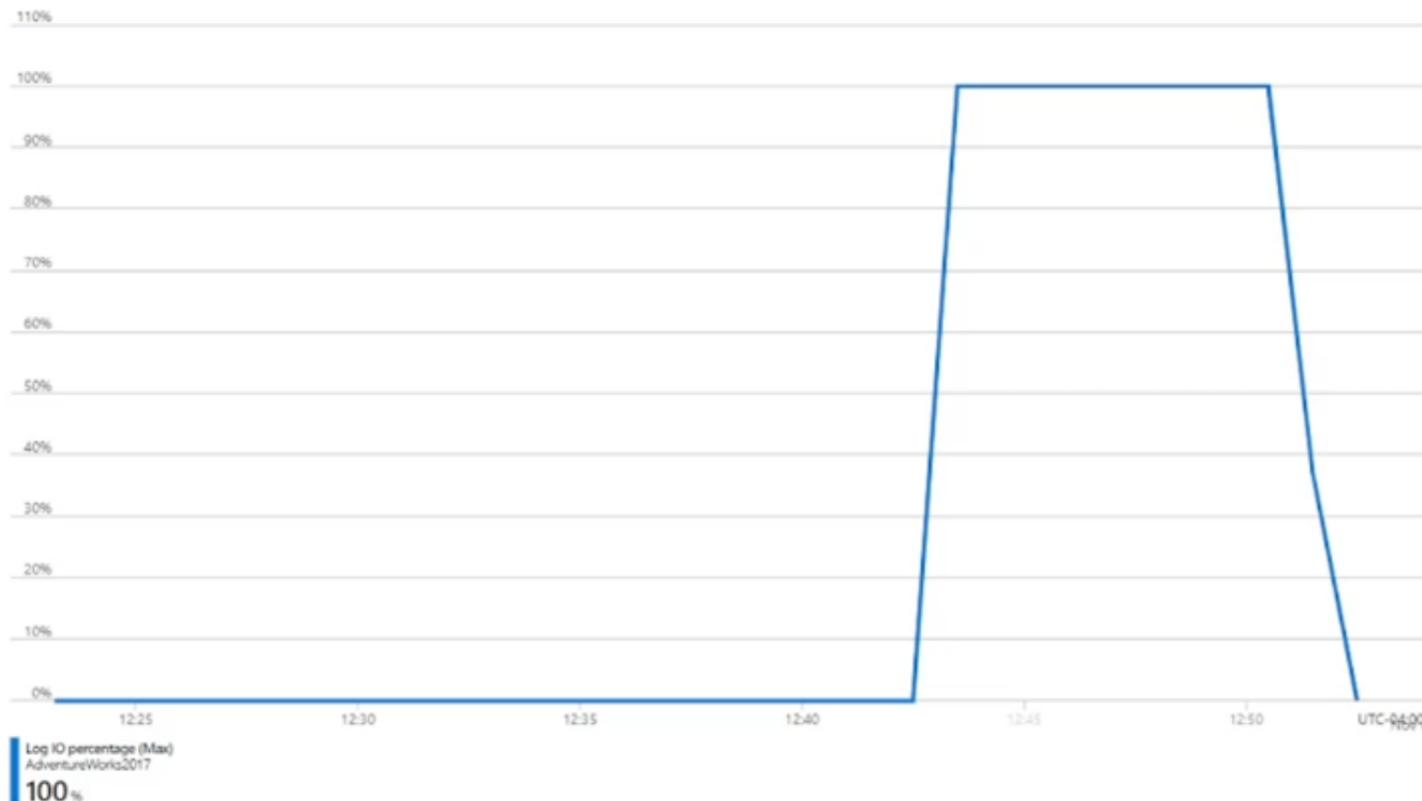
Reference:

<https://databricks.com/blog/2020/03/25/trust-but-verify-with-databricks.html>

NEW QUESTION 42

- (Exam Topic 5)

You have an Azure SQL database named DB1 in the General Purpose service tier. The performance metrics for DB1 are shown in the following exhibit.



You need to reduce the Log 10 percentage. The solution must minimize costs. What should you do?

- A. Increase the number of vCores.
- B. Change RecoverymodeltoSimple.
- C. Performcheckpoint operation.
- D. ChangeService tier toBusiness Critical.

Answer: D

NEW QUESTION 43

- (Exam Topic 5)

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive
- C. High Concurrency

Answer: A

Explanation:

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.

The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

NEW QUESTION 46

- (Exam Topic 5)

You have an Azure SQL database named db1.

You need to retrieve the resource usage of db1 from the last week.

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

NOTE:Each correct selection is worth one point.

SELECT *

FROM

| |
|------------------------------------|
| ▼ |
| sys.dm_db_resource_stats |
| sys.dm_exec_requests |
| sys.dm_user_db_resource_governance |
| sys.resource_stats |

WHERE database_name = 'db1' AND

start_time >

| |
|------------------|
| ▼ |
| DATEADD |
| DATEDIFF |
| DATEPART |
| TODATETIMEOFFSET |

(day, -7, GETDATE())

ORDER BY start_time DESC;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.resource_stats

sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.

Box 2: DateAdd

The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

DECLARE @s datetime; DECLARE @e datetime;

SET @s= DateAdd(d,-7,GetUTCDate()); SET @e= GETUTCDATE();

SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats

WHERE start_time BETWEEN @s AND @e GROUP BY database_name

HAVING AVG(avg_cpu_percent) >= 80 Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-data>

NEW QUESTION 48

- (Exam Topic 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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.

You need to improve system performance. Solution: You create additional tempdb files. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Reference:
<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION 51

- (Exam Topic 5)
 You have the following Transact-SQL query.

```
SELECT
    [file_id] AS [File ID],
    [type] AS [File Type],
    substring([physical_name], 1,1) AS [Drive],
    [name] AS [Logical Name],
    [physical_name] AS [Physical Name],
    CAST([size] as DECIMAL(38,0))/128.0 AS [ColumnA],
    CAST(FILEPROPERTY([name], 'SpaceUsed') AS DECIMAL(38,0))/128.0 AS
[ColumnB],
    (CAST([size] AS DECIMAL(38,0))/128.0) - (CAST(FILEPROPERTY([name],
'SpaceUsed') AS DECIMAL (38,0))/128.0) AS [ColumnC],
    [max_size] AS [ColumnD],
    [is_percent_growth] AS [Percent Growth Enabled],
    [growth] AS [Growth Rate],
    SYSDATETIME() AS [Current Date]
FROM sys.database_files;
```

Which column returned by the query represents the free space in each file?

- A. ColumnA
- B. ColumnB
- C. ColumnC
- D. ColumnD

Answer: C

Explanation:

Example:
 Free space for the file in the below query result set will be returned by the FreeSpaceMB column. SELECT DB_NAME() AS DbName, name AS FileName, type_desc, size/128.0 AS CurrentSizeMB, size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB FROM sys.database_files WHERE type IN (0,1);
 Reference:
<https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/>

NEW QUESTION 55

- (Exam Topic 5)
 You have an Azure SQL managed instance.
 You need to restore a database named DB1 by using Transact-SQL.
 Which command should you run? To answer, select the appropriate options in the answer area.
 NOTE:Each correct selection is worth one point.

RESTORE DB1 FROM

| | |
|--|--|
| <input type="text" value="DATABASE"/> <input type="text" value="FILE"/> <input type="text" value="LOG"/> | <input type="text" value="DISK = N'\\NAS01\SSQLBackups\DB1.bak';"/> <input type="text" value="TAPE = N'\\Tape0'"/> <input type="text" value="URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'"/> |
|--|--|

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Text Description automatically generated

NEW QUESTION 56

- (Exam Topic 5)
 You have an Azure SQL database named DB1 in the General Purpose service tier. You need to monitor D81 by using SQL insights.
 What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

To collect monitoring data, use:

| | |
|---|---|
| | ▼ |
| A virtual machine An Azure function The Azure Monitor agent | |

To store monitoring data, create:

| | |
|--|---|
| | ▼ |
| A Log Analytics workspace An Azure SQL database An Azure Storage account | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To collect monitoring data, use:

| | |
|---|---|
| | ▼ |
| A virtual machine An Azure function The Azure Monitor agent | |

To store monitoring data, create:

| | |
|--|---|
| | ▼ |
| A Log Analytics workspace An Azure SQL database An Azure Storage account | |

NEW QUESTION 59

- (Exam Topic 5)

You are planning disaster recovery for the failover group of an Azure SQL Database managed instance. Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs. You set the Read/Write failover policy to Automatic. What are two results of the configuration? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. In the event of a datacenter or Azure regional outage, the databases will fail over automatically.
- B. In the event of an outage, the databases in the primary instance will fail over immediately.
- C. In the event of an outage, you can selectively fail over individual databases.
- D. In the event of an outage, you can set a different grace period to fail over each database.
- E. In the event of an outage, the minimum delay for the databases to fail over in the primary instance will be one hour.

Answer: AE

Explanation:

A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.
 E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.

Reference:

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

NEW QUESTION 63

- (Exam Topic 5)

You have An Azure SQL managed instance. You need to configure the SQL Server Agent service to email job notifications. Which statement should you execute?

A)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
```

B)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'application_dbmail_profile';
```

C)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
```

D)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sys_dbmail_profile';
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

NEW QUESTION 67

- (Exam Topic 5)

You have an Azure subscription that contains a SQL Server on Azure Virtual Machines instance named SQLVMI. SQLVMI hosts a database named OBI. You need to retrieve query plans from the Query Store on DB1. What should you do first?

- A. On SQLVMI1, install the SQL Server IaaS Agent extension.
- B. From Microsoft SQL Server Management Studio, modify the properties of the SQL Server instance.
- C. From Microsoft SQL Server Management Studio, modify the properties of DB 1.
- D. On SQLVMI1, install the Azure Monitor agent for Windows.

Answer: B

NEW QUESTION 72

- (Exam Topic 5)

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service. You have an Azure Key vault named vault1 that contains an encryption key named key1. You need to encrypt df1 by using key1. What should you do first?

- A. Disable purge protection on vault1.
- B. Remove the linked service from df1.
- C. Create a self-hosted integration runtime.
- D. Disable soft delete on vault1.

Answer: B

Explanation:

A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.

Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key>

NEW QUESTION 77

- (Exam Topic 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 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

Answer: A

Explanation:

The REPLACE option overrides several important safety checks that restore normally performs. The overridden checks are as follows:

- Restoring over an existing database with a backup taken of another database.

With the REPLACE option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

NEW QUESTION 79

- (Exam Topic 5)

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

| Name | Type |
|------|---|
| App1 | Azure web app |
| db1 | Azure SQL database in the serverless tier |

App1 experiences transient connection errors and timeouts when it attempts to access db1 after extended periods of inactivity. You need to modify db1 to resolve the issues experienced by App1 as soon as possible, without considering immediate costs. What should you do?

- A. Increase the number Of vCores allocated to db1.
- B. Disable auto-pause delay for db1.
- C. Decrease the auto-pause delay for db1.

D. Enable automatic tuning for db1.

Answer: D

NEW QUESTION 80

- (Exam Topic 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

Answer: A

Explanation:

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>

NEW QUESTION 85

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1.

What should you do?

- A. Connect to db1 and run the sp_configure command.
- B. Connect to the master database of server1 and run the sp_configure command.
- C. Configure the extended properties of db1.
- D. Modify the database scoped configuration of db1.

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism>

NEW QUESTION 88

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Hyperscale
- B. Azure SQL Database managed instance General Purpose
- C. Azure SQL Database Premium
- D. Azure SQL Database Basic

Answer: C

Explanation:

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). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.

Reference:

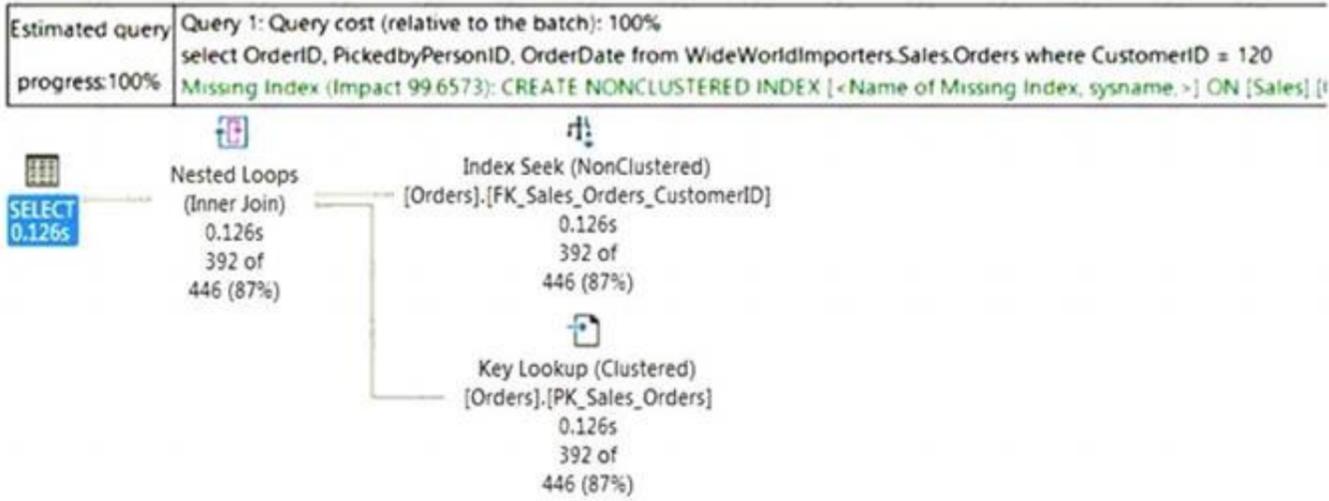
<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

NEW QUESTION 89

- (Exam Topic 5)

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.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

Reference:

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

NEW QUESTION 94

- (Exam Topic 5)

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:

The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.

Administrative effort must be minimized. What should you use to host the databases?

- A. a single Azure SQL database
- B. an Azure SQL Database elastic pool
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed Instance

Answer: D

NEW QUESTION 98

- (Exam Topic 5)

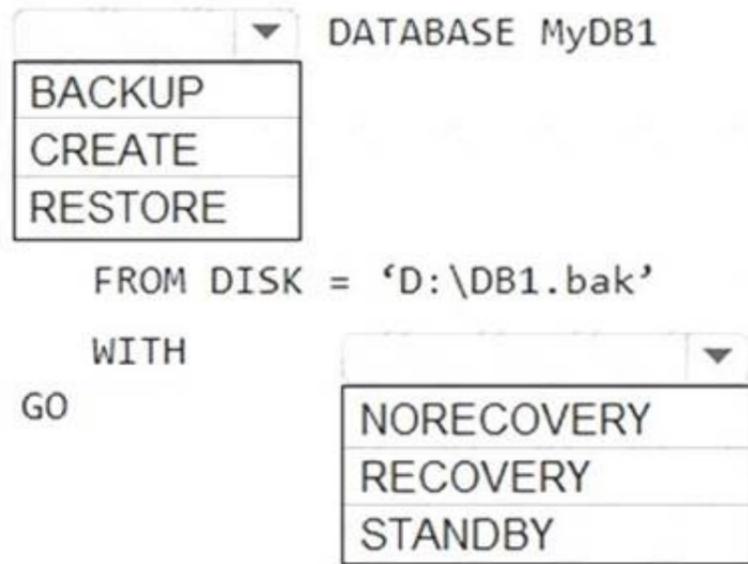
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.

You plan to deploy an Always On availability group that will have the following configurations:

- > VM1 will host the primary replica of DB1.
- > VM2 will host a secondary replica of DB1.

You need to prepare the secondary database on VM2 for the availability group.

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondar>

NEW QUESTION 99

- (Exam Topic 5)

You are building a database backup solution for a SQL Server database hosted on an Azure virtual machine. In the event of an Azure regional outage, you need to be able to restore the database backups. The solution must minimize costs.

Which type of storage accounts should you use for the backups?

- A. locally-redundant storage (LRS)
- B. read-access geo-redundant storage (RA-GRS)
- C. zone-redundant storage (ZRS)
- D. geo-redundant storage

Answer: B

Explanation:

Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages. However, that data is available to be read only if the customer or Microsoft initiates a failover from the primary to secondary region. When you enable read access to the secondary region, your data is available to be read if the primary region becomes unavailable. For read access to the secondary region, enable read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>

NEW QUESTION 104

- (Exam Topic 5)

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1. You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

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

NOTE:Each correct selection is worth one point.

- A. From the properties of DB1. configure geo-replication.
- B. From the properties of Server1 add a failover group.
- C. Create a new Azure SQL Database server named Server2.
- D. From the properties of Server1 configure retention for DB1
- E. Create a new Azure SQL Database instance named DB2.

Answer: BC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell> <https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azur>

NEW QUESTION 107

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 contains a table named CustomerPII.

You need to record whenever users query the CustomerPII table.

Which two options should you enable? Each correct answer presents part of the solution.

NOTE:Each correct selection is worth one point.

- A. server audit specification
- B. SQL Server audit
- C. database audit specification
- D. a server principal

Answer: AC

Explanation:

An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):

- A server policy applies to all existing and newly created databases on the server.
- If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings.
- Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.

Note:

The Server Audit Specification object belongs to an audit.

A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auditing-overview>

NEW QUESTION 110

- (Exam Topic 5)

You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.

You need to ensure that the design provides the fastest query times of the records when querying for arbitrary date ranges and aggregating by fiscal calendar attributes.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Create a date dimension table that has a DateTime key.
- B. Create a date dimension table that has an integer key in the format of YYYYMMDD.
- C. Use built-in SQL functions to extract date attributes.
- D. Use integer columns for the date fields.
- E. Use DateTime columns for the date fields.

Answer: BD

Explanation:

Reference:

https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-ina

NEW QUESTION 111

- (Exam Topic 5)

You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases
- B. geo-replication support
- C. number of concurrently peaking databases * peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases * average CPU utilization per database

Answer: ACE

Explanation:

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model: $\text{MAX}(\langle \text{Total number of DBs} \times \text{average vCore utilization per DB} \rangle, \langle \text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB} \rangle)$

A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

NEW QUESTION 113

- (Exam Topic 5)

You have an Azure subscription that contain an Azure SQL managed instance named SQLMI1 and a Log Analytics workspace named Workspace1.

You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace1.

- A. Create the private endpoint connection on SQLMI1.
- B. Configure Azure SQL Analytics to use Workspace1.
- C. Modify the Computer + storage settings for SQLMI1.
- D. Modify the diagnostic settings for SQLMI1.

Answer: B

NEW QUESTION 116

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

Which two actions should you perform? Each correct answer presents part of the solution.
 NOTE: Each correct selection is worth one point.

- A. EnableLast_Query_Plan_Statsin the master database
- B. EnableLightweight_Query_Profilingin DB1
- C. EnableLast_Query_Plan_Statsin DB1
- D. EnableLightweight_Query_Profilingin the master database
- E. EnablePARAMETER_SNIFFINGin DB1

Answer: AC

Explanation:

Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm_exec_query_plan_stats.

Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

NEW QUESTION 118

- (Exam Topic 5)

You have an Azure subscription.

You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.

You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs. The solution must minimize administrative effort.

Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

MSX:

| | |
|-----------------------|---|
| | ▼ |
| SQL database | |
| SQL managed instances | |
| SQL virtual machines | |

TSX:

| | |
|-----------------------|---|
| | ▼ |
| SQL database | |
| SQL managed instances | |
| SQL virtual machines | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

MSX: ▼
 SQL database
 SQL managed instances
 SQL virtual machines

TSX: ▼
 SQL database
 SQL managed instances
 SQL virtual machines

NEW QUESTION 120

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.

You need to resolve the performance issues. What should you do?

- A. Implement memory-optimized tables.
- B. Run the dbcc flushprocindbcommand.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the dbcc dbreindexcommand.

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb>

NEW QUESTION 123

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy a logical SQL server by using PowerShell. The solution must ensure that the logical SQL server can create Azure AD users and provide Transparent Data Encryption (TDE) with a customer-managed key.

How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```
New-AzSqlServer -ResourceGroupName "RG1" -ServerName "SQL1" -Location "EastUS" -ErrorAction Stop
-Tags @{Environment="Databases";Department="Data Tech"}
-assignidentity
-federatedclientID
-keyid
-assignidentity
-federatedclientID
-keyid
"https://db1.vault.azure.net/keys/dbkey/01234234512345678901234561823942"
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

"Assigned Identity" and "Key

id" <https://learn.microsoft.com/en-us/powershell/module/az.sql/new-azsqlserver?view=azps-10.2.0#code-try-3>

NEW QUESTION 124

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains a private certificate named Sales. The private key for Sales is encrypted with a password. You need to change the password for the private key. Which

Transact-SQL statement should you run?

```
A)
ALTER CERTIFICATE Sales
WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%',
ENCRIPTION BY PASSWORD = '6YY9Ycd!pV');
```

B)
 ALTER CERTIFICATE Sales
 WITH PRIVATE KEY (ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');

C)
 ALTER CERTIFICATE Sales WITH PRIVATE KEY (FILE = 'D:\importkeys\SalesNew, DECRYPTION BY PASSWORD = ' Mb^6BK&*w%');

D)
 ALTER CERTIFICATE Sales WITH PRIVATE KEY (DECRYPTION BY PASSWORD = ' EWYx9Xk+\${#}');

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

NEW QUESTION 126

- (Exam Topic 5)

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.

Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

MFA:

- Azure AD authentication
- Microsoft SQL Server authentication
- Passwordless authentication
- Windows authentication

Database-level authentication:

- Application roles
- Contained database users
- Database roles
- Microsoft SQL Server logins

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated

Box 1: Azure AD authentication

Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.

Box 2: Contained database users

Azure Active Directory Uses contained database users to authenticate identities at the database level. Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-authentication>

NEW QUESTION 130

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:

- Minimize how long it takes to fail over.
- Maintain existing connections to the primary replica during a failover. What should you do?

- A. Connect each virtual machine to a single subnet on a single virtual network.
- B. Connect each virtual machine to a single subnet on a virtual network
- C. Deploy a standard Azure load balancer.
- D. Connect each virtual machine to a different subnet on a single virtual network.
- E. Connect each virtual machine to a different subnet on a virtual network
- F. Deploy a basic Azure load balancer.

Answer: C

NEW QUESTION 131

- (Exam Topic 5)

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- > Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
- > Minimize the possibility of VM1 exfiltrating data stored in SqlDb1. What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

Answer: C

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

NEW QUESTION 136

- (Exam Topic 5)

You have an Azure SQL database.

You need to identify whether a delayed query execution is associated to a RESOURCE wait.

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

NOTE: Each correct selection is worth one point.

Answer Area

```

SELECT [wait_type]
      [wait_type]
      context_info
      wait_resource

SUM(wait_time) AS total_wait_time_ms

FROM [sys.dm_exec_requests] AS dm1
JOIN [sys.dm_exec_connections]
     [sys.dm_db_partition_stats] ON dm1.connection_id

WHERE is_user_process = 1

GROUP BY TARGET1

ORDER BY SUM(wait_time) DESC;
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

SELECT [wait_type]
      [wait_type]
      context_info
      wait_resource

SUM(wait_time) AS total_wait_time_ms

FROM [sys.dm_exec_requests] AS dm1
JOIN [sys.dm_exec_connections]
     [sys.dm_db_partition_stats] ON dm1.connection_id

WHERE is_user_process = 1

GROUP BY TARGET1

ORDER BY SUM(wait_time) DESC;
    
```

NEW QUESTION 139

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.

How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Encryption key: ▼

| |
|-------------------------|
| Column encryption key |
| Database encryption key |
| Service master key |

Encryption type: ▼

| |
|-----------------------------------|
| Deterministic |
| Randomized |
| Transparent Data Encryption (TDE) |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION 141

- (Exam Topic 5)

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: > Minimize costs.

> 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.

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo>

NEW QUESTION 142

- (Exam Topic 5)

You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1. You need to ensure that User1 can create databases on Server1. Which database role should you assign to User1?

- A. db_owner
- B. dbmanager
- C. dbo
- D. db_ddladmin

Answer: B

NEW QUESTION 143

- (Exam Topic 5)

You have an Azure SQL Database elastic pool that contains 10 databases. You receive the following alert.

Msg 1132, Level 16, State 1, Line 1

The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.

You need to resolve the alert. The solution must minimize administrative effort.

Which three actions can you perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Delete data from a database.
- B. Remove a database from the pool.
- C. Increase the maximum storage of the elastic pool.
- D. Shrink individual databases.
- E. Enable data compression.

Answer: BCD

NEW QUESTION 146

- (Exam Topic 5)

You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information. You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information. What should you include in the recommendation?

- A. row-level security
- B. data masking
- C. Always Encrypted
- D. column-level security

Answer: B

Explanation:

Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users. The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card. Example: XXXX-XXXX-XXXX-1234

NEW QUESTION 149

- (Exam Topic 5)

You have SQL Server 2019 or an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory. You scale up the virtual machine to 8 vCPUs and 64 GB of memory. You need to reduce tempdb contention without negatively affecting server performance. What is the number of secondary data files that you should configure for tempdb?

- A. 2
- B. 4
- C. 8
- D. 64

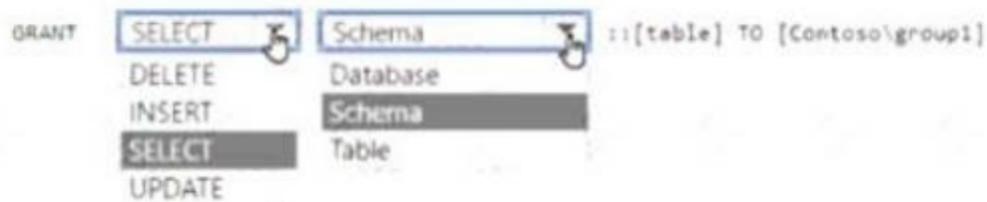
Answer: C

NEW QUESTION 151

- (Exam Topic 5)

You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created in 081. 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.

Answer Area

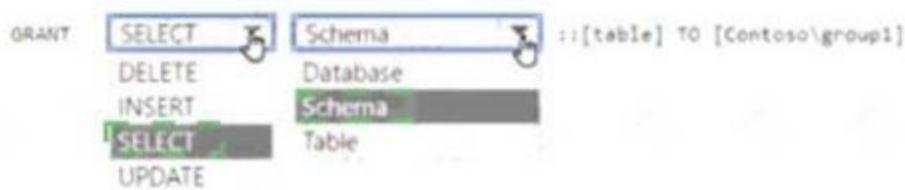


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 153

- (Exam Topic 5)

You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently. You enable columnstore archival compression. What are two possible results of the configuration? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Queries that use the index will consume more disk I/O.
- B. Queries that use the index will retrieve fewer data pages.
- C. The index will consume more disk space.
- D. The index will consume more memory.
- E. Queries that use the index will consume more CPU resources.

Answer: BE

Explanation:

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk. Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

NEW QUESTION 155

- (Exam Topic 5)

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.

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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

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

NEW QUESTION 160

- (Exam Topic 5)

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- The web app must be hosted on an Azure virtual network.
 - The Azure SQL database must be assigned a private IP address.
 - The Azure SQL database must allow connections only from the virtual network. You need to recommend a solution that meets the requirements.
- What should you include in the recommendation?

- A. Azure Private Link
- B. a network security group (NSG)
- C. a database-level firewall
- D. a server-level firewall

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/private-endpoint-overview>

NEW QUESTION 163

- (Exam Topic 5)

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.

You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

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

NOTE:Each correct selection is worth one point.

Distribution: ▼

- Hash
- Round robin
- Replicated

Index: ▼

- Clustered columnstore
- Clustered
- Nonclustered

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, table, chat or text message Description automatically generated

Box 1: Hash

Consider using a hash-distributed table when: The table size on disk is more than 2 GB.

The table has frequent insert, update, and delete operations. Box 2: Clustered columnstore

Clustered columnstore tables offer both the highest level of data compression and the best overall query performance.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-index>

NEW QUESTION 164

- (Exam Topic 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 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: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 169

- (Exam Topic 5)

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

Answer: AD

NEW QUESTION 174

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table. You need to identify the extent of the data skew in Table1. What should you do in Synapse Studio?

- A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
- B. Connect to the built-in pool and run DBCC CHECKALLOC.
- C. Connect to Pool1 and run DBCC CHECKALLOC.
- D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Answer: D

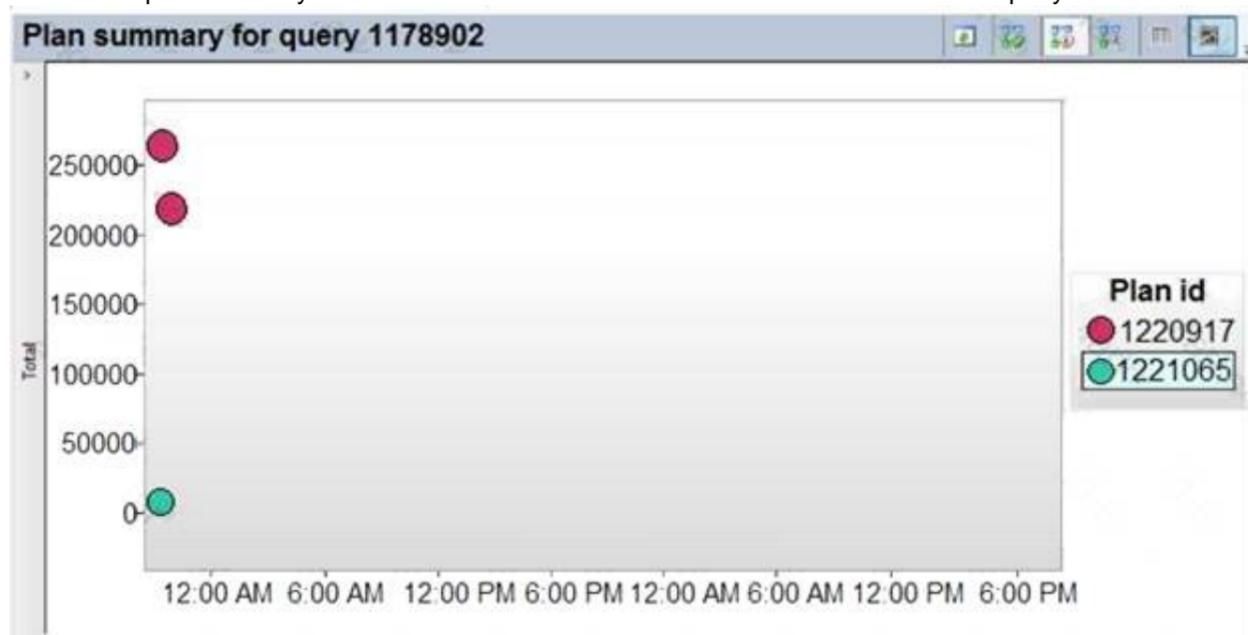
Explanation:

Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data. Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet>

NEW QUESTION 175

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the DBCC FREEPROCCACHE command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

Answer: C

Explanation:

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

NEW QUESTION 177

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named App1. App1 uses db1. You need to enable Resource Governor for a App1. The solution must meet the following requirements: App1 must be able to consume all available CPU resources.

App1 must have at least half of the available CPU resources always available.

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

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

- Create a plan.
- Create a classifier function in db1.
- Create a workload group.
- Create a classifier function in the master database.
- Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

➔

➜

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 181

- (Exam Topic 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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.

You need to improve system performance.

Solution: You change the data file for the master database to autogrow by 10 percent. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION 183

- (Exam Topic 5)

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1. You need to configure SQL1 to use mixed mode authentication. Which procedure should you run?

- A. sp_addremotelogin
- B. xp_instance_regwrite
- C. sp_cnchange_users_login
- D. xp_grant_login

Answer: B

NEW QUESTION 188

- (Exam Topic 5)

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

Answer: B

Explanation:

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>

NEW QUESTION 191

- (Exam Topic 5)

You are building an Azure Stream Analytics job to retrieve game data.

You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.

How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
SELECT  as HighestScore
FROM input TIMESTAMP BY CreatedAt
GROUP BY 
        CollectTop(1)OVER(ORDER BY Score Desc)
        Game, MAX(Score)
        TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.

Analytic Function Syntax:

TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT

DURATION clause> [<WHEN clause>])

Box 2: Tumbling(minute 5)

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

Tell me the count of Tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Reference:

<https://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun>

NEW QUESTION 196

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. distributed availability groups
- B. database mirroring
- C. log shipping
- D. Database Migration Assistant

Answer: D

Explanation:

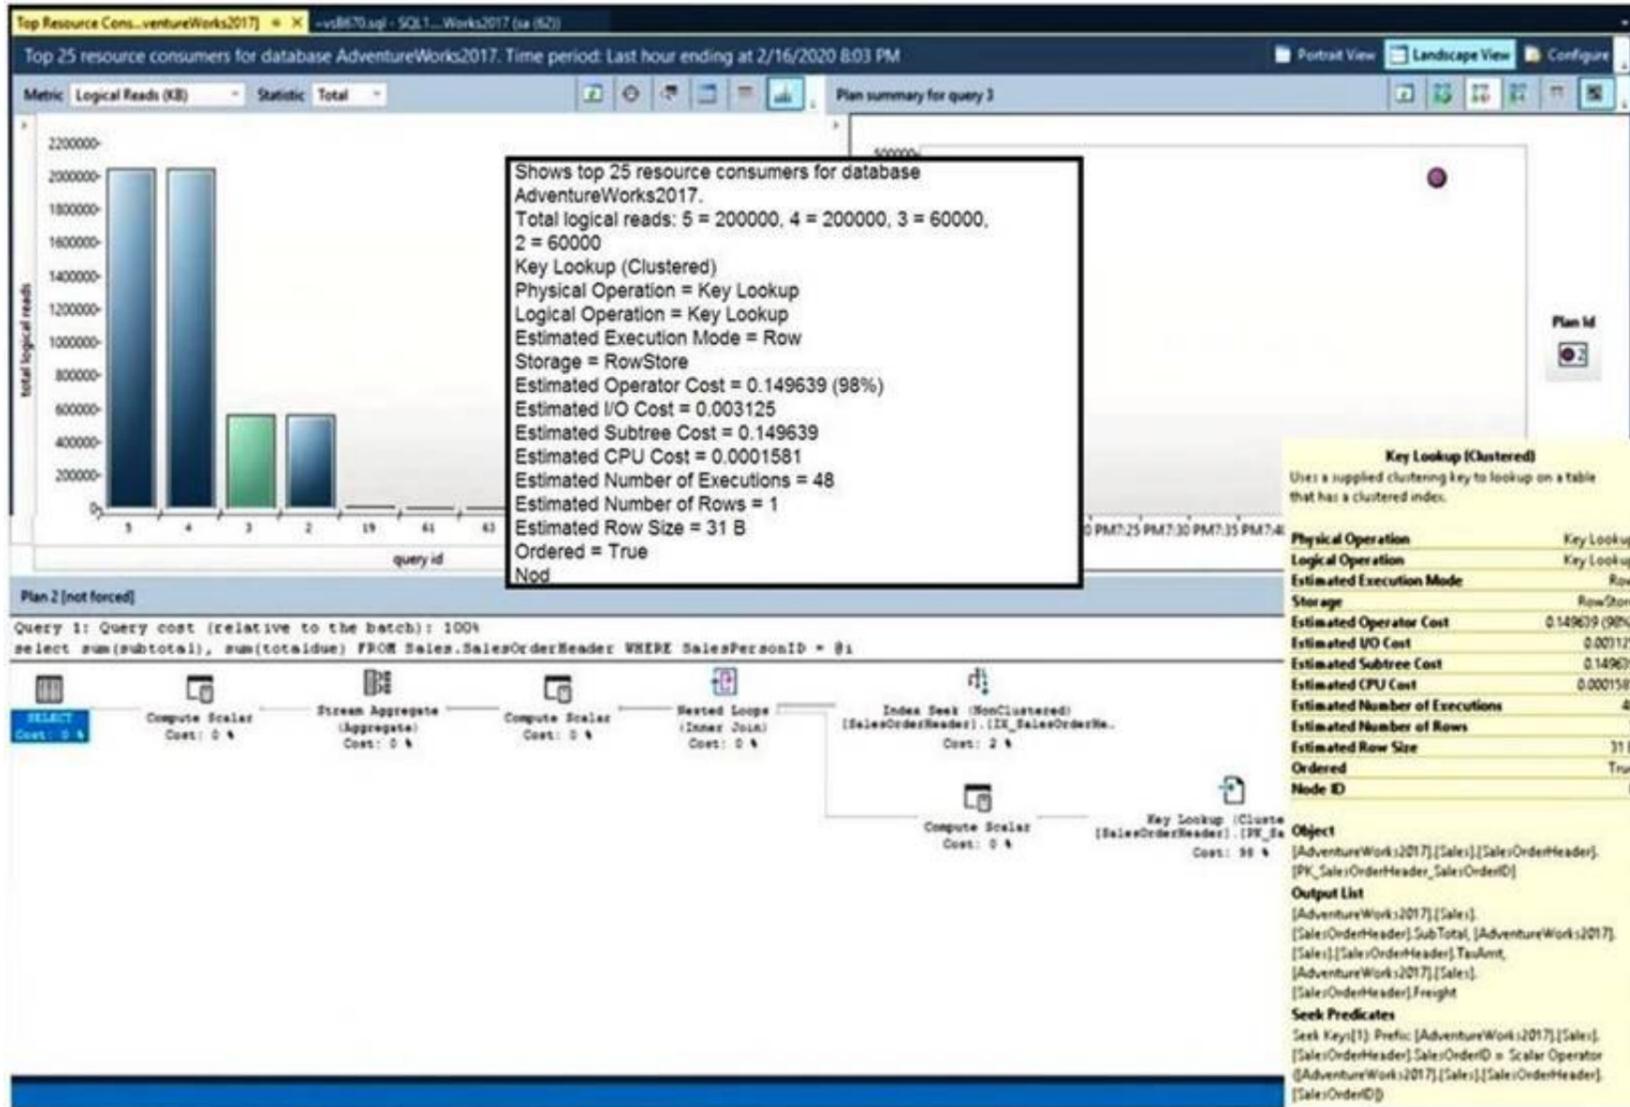
Ref:<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-azure-sql>

NEW QUESTION 200

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



For each of the following statements, select yes if the statement is true. Otherwise, select no.

NOTE: Each correct selection is worth one point.

Statements

You will reduce the I/O usage and the query execution time if you force the query plan.

Yes

No

You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.

You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu>

NEW QUESTION 204

- (Exam Topic 5)

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
    payment_type,
    SUM(fare_amount) AS fare_total
FROM OPENROWSET (
    BULK 'csv/busfare/tripdata_2020*.csv',
    DATA_SOURCE = 'BusData',
    FORMAT = 'CSV', PARSER_VERSION = '2.0',
    FIRSTROW = 2
)
WITH (
    payment_type INT 10,
    fare_amount FLOAT 11
) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

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

The query results include only [answer choice] in the csv/busfare folder.

▼

CSV files in the tripdata_2020 subfolder
 files that have files names beginning with "tripdata_2020"
 CSV files that have file names containing "tripdata_202"
 CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is [answer choice] row.

▼

a header
 a data
 an empty

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, table Description automatically generated
 Box 1: CSV files that have file named beginning with "tripdata_2020" Box 2: a header
 FIRSTROW = 'first_row'
 Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.
 Reference:
<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset>

NEW QUESTION 208

- (Exam Topic 5)

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary. You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions | | Answer Area | |
|---|---|-------------|---|
| Encrypt the Salary column by using the randomized encryption type. | | | |
| Create a column encryption key. | | | |
| Enable Transparent Data Encryption (TDE). | ➤ | | ⬆ |
| Encrypt the Salary column by using the deterministic encryption type. | ⬅ | | ⬇ |
| Apply a dynamic data mask to the Salary column. | | | |
| Create a column master key. | | | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION 212

- (Exam Topic 5)

You configure backup for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore
 Specify how long you want to keep your point-in-time backups. [Learn more >](#)

How many days would you like PITR backups to be kept?

Long-term retention
 Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more >](#)

Weekly LTR Backups
 Keep weekly backups for:

Monthly LTR Backups
 Keep the first backup of each month for:

Yearly LTR Backups
 Keep an annual backup for:

Which weekly backup of the year would you like to keep?

Use the drop-down menus to select the answer choice the completes each statement based on the information presented in the graphic.
 NOTE: Each correct selection is worth one point.

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**

- point-time restore (PITR) backup.
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.

- 65 backup copies
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

- point-time restore (PITR) backup
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

- 65 backup copies
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

NEW QUESTION 213

- (Exam Topic 5)

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server. You need to ensure that each customer can only connect to and access their respective database. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.
- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

Answer: BE

Explanation:

Manage database access by adding users to the database, or allowing user access with secure connection strings. Database-level firewall rules only apply to individual databases. Reference: <https://docs.microsoft.com/en-us/azure/sql/database/secure-database-tutorial>

NEW QUESTION 215

- (Exam Topic 5)

You have an Azure subscription. You plan to deploy an Azure SQL database by using an Azure Resource Manager template. How should you complete the template? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "resources": [
    {
      "type": 
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",
          ...
          "dependsOn": [
            "properties": [
              "tags": [
                "[resourceId('Microsoft.Sql/servers', concat(parameters('name1')))]"
              ]
            ]
          ]
        }
      ]
    }
  ]
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart>

NEW QUESTION 217

- (Exam Topic 5)

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE:Each correct selection is worth one point.

Values

Answer Area

- CustomerKey
- HASH
- ROUND_ROBIN
- REPLICATE
- OrderDateKey
- SalesOrderNumber

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [ ] ([ProductKey])
, PARTITION ( [ ] ] RANGE RIGHT FOR VALUES
(20170101, 20180101, 20190101, 20200101, 20210101)
)
)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: HASH

Box 2: OrderDateKey

In most cases, table partitions are created on a date column.

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool>

NEW QUESTION 222

- (Exam Topic 5)

You have an Azure Synapse Analytics Apache Spark pool named Pool1.

You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.

You need to load the files into the tables. The solution must maintain the source data types. What should you do?

- A. Load the data by using PySpark.
- B. Load the data by using the OPENROWSET Transact-SQL command in an Azure Synapse Analytics serverless SQL pool.
- C. Use a Get Metadata activity in Azure Data Factory.
- D. Use a Conditional Split transformation in an Azure Synapse data flow.

Answer: B

Explanation:

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.

To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.

The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files> <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage>

NEW QUESTION 225

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You need to display the estimated execution plan of a query by using the query editor in the Azure portal.

What should you do first?

- A. Run the `showplan_all` Transact-SQL statement.
- B. For DB1, set `QUERY_CAPTURE_MODE` of Query Store to All.
- C. Run the `forceplan` Transact-SQL statement.
- D. Enable Query Store for DB1.

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-showplan-all-transact-sql?view=sql-server-ver15>

NEW QUESTION 228

- (Exam Topic 5)

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning

Inherit from:

Server
Azure defaults
Don't inherit

The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options

| OPTION | DESIRED STATE | CURRENT STATE |
|--------------|---|---------------------------------------|
| FORCE PLAN | <input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT | ON Auto-configured by Azure |
| CREATE INDEX | <input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT | ON Auto-configured by Azure |
| DROP INDEX | <input checked="" type="radio"/> ON <input type="radio"/> OFF <input type="radio"/> INHERIT | ON Forced by user |

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

| Statements | Yes | No |
|--|-----------------------|-----------------------|
| Nonclustered indexes will be added to tables to improve performance. | <input type="radio"/> | <input type="radio"/> |
| Columns will be added to existing indexes automatically. | <input type="radio"/> | <input type="radio"/> |
| The query execution plan will revert to a previous plan if query performance degrades. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

We see: Tuning option: Create index ON

CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.

Box 2: No

Box 3: Yes

FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

NEW QUESTION 233

- (Exam Topic 5)

You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2. Which statement should you run?

- A. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);
- B. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (COPY_ONLY, STATS=5, COMPRESSION);
- C. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (File_Snapshot, STATS=5, COMPRESSION);
- D. BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/ mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);

Answer: B

Explanation:

BACKUP DATABASE supports only copy-only full backups of databases, files, or filegroups when it's executed on secondary replicas. Copy-only backups don't impact the log chain or clear the differential bitmap.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on>

NEW QUESTION 238

- (Exam Topic 5)

You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions | Answer Area |
|--|-------------|
| Create a workload group. | |
| Create a user-defined classifier function. | ⬅️ ⬆️ |
| Modify Resource Governor. | ➡️ ⬆️ |
| Create a contained database user. | |
| Create a resource pool. | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server> <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-def>

NEW QUESTION 241

- (Exam Topic 5)

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

- A. role assignments
- B. account keys
- C. shared access signatures (SAS)
- D. Azure Active Directory (Azure AD) identities

Answer: C

Explanation:

A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:

What resources the client may access.

What permissions they have to those resources. How long the SAS is valid.

Note: Data Lake Storage Gen2 supports the following authorization mechanisms:

- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Shared Key authorization
- > Shared access signature (SAS) authorization
- > Role-based access control (Azure RBAC)
- > Access control lists (ACL)

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

NEW QUESTION 245

- (Exam Topic 5)

You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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

| Actions | Answer Area |
|--|-------------|
| Duplicate Azure Policy definitions | |
| Run Azure Policy remediation tasks | |
| Create an Azure Blueprints assignment | ⬆ |
| Create an Azure Policy initiative | ⬇ |
| Create an Azure Policy initiative assignment | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create an Azure Policy Initiative

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.

With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

Step 2: Create an Azure Policy Initiative assignment

Assign the initiative definition you created in the previous step. Step 3: Run Azure Policy remediation tasks

To apply the Policy Initiative to the existing SQL databases. Reference:

<https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

NEW QUESTION 247

- (Exam Topic 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 SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query `sys.dm_exec_requests` and discover that the wait type is `PAGELATCH_UP` and the `wait_resource` is `2:3:905856`.

You need to improve system performance. Solution: You shrink the transaction log file. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION 252

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template. How should you complete the ARM template? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

Answer Area

```

"resources": [
  ...
  "type": [
    Microsoft.Sql/servers
    Microsoft.Sql/servers/databases
    Microsoft.Sql/managedInstances
  ],
  "name": "[parameters('targetName')]",
  "location": "[parameters('location')]",
  "sku": {
    "name": "[parameters('skuName')]"
  },
  ...
  "dependsOn": [
    "[parameters('targetName')]",
    "[parameters('virtualNetworkName')]",
    "[variables('networkSecurityGroupName')]",
  ],
  "properties": {
    "administratorLogin": "[parameters('administratorLogin')]",
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('virtualNetworkName'), parameters('virtualNetworkName'), parameters('subnetName'))]",
    "storageSizeInGB": "[parameters('storageSizeInGB')]", "vCores": "[parameters('vCores')]",
    "licenseType": "[parameters('licenseType')]"
  },
  ...
]

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, Word, email Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/create-template-quickstart?tabs=azure-powe>

NEW QUESTION 253

- (Exam Topic 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 SQL database named Sales.

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

- > During normal operations, provide at least two readable copies of Sales.
- > Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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>

NEW QUESTION 256

- (Exam Topic 5)

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 inDBV 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.ddladmm role.
- B. Assign the database user the db.owner role.
- C. Create a contained database user.
- D. Create a login and an associated database user.

Answer: AD

NEW QUESTION 261

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an instance of SQL Sevier on Azure Virtual Machines that supports Write Accelerator.

Which virtual machine series should you use?

- A. H-series
- B. G -series
- C. M-series
- D. E-series

Answer: C

NEW QUESTION 262

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine named SQL1. SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator. You need to ensure that dbadmin1 receives an email alert if a job fails.

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

| Actions | | Answer Area |
|--|---|-------------|
| Create a job alert | | |
| Create a job notification | | |
| Enable Database Mail | ➤ | ⬆ |
| Enable the email settings for the SQL Server Agent | ⬅ | ⬇ |
| Create a job target | | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Enable the email settings for the SQL Server Agent.

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail. Step 2: Create a job alert

Step 3: Create a job notification Example:

```
-- adds an e-mail notification for the specified alert (Test Alert)
-- This example assumes that Test Alert already exists
-- and that François Ajenstat is a valid operator name. USE msdb ;
GO
```

```
EXEC dbo.sp_add_notification
@alert_name = N'Test Alert',
@operator_name = N'François Ajenstat',
@notification_method = 1 ; GO
```

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status> <https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator>

NEW QUESTION 263

- (Exam Topic 5)

Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

- A. Hyperscale
- B. General Purpose
- C. Premium
- D. Business Critical

Answer: D

NEW QUESTION 267

- (Exam Topic 5)

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.

Answer: B

NEW QUESTION 270

- (Exam Topic 5)

You have an Azure SQL managed instance.

You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries.

What should you do?

- A. Generate an estimated execution plan.
- B. Generate an actual execution plan.
- C. Run sys.dm_exec_query_plan_stats.
- D. Generate Live Query Statistics.

Answer: C

Explanation:

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-quer>

NEW QUESTION 275

- (Exam Topic 5)

You need to use an Azure Resource Manager (ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files

How should you complete the template? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,
    ...
  }
}

"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {"name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]"},
      "input": {"lun": "[copyIndex('dataDisks')],"createOption": "empty",
      "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),parameters('dataDiskCount')),
        variables('dataDisks').caching )]", "diskSizeGB": 1023,
  
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Read onlyReadWrite

NEW QUESTION 276

- (Exam Topic 5)

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

- Support real-time data replication to a different geographic region.
- Use Azure as a disaster recovery target.
- 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

Answer: D

NEW QUESTION 278

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named DB1. The database reports a CHECKSUM error.

You need to recover the database.

How should you complete the statements? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

USE master;

ALTER DATABASE [DB1] SET

| | |
|-------------|---|
| | ▼ |
| OFFLINE | |
| ONLINE | |
| SINGLE_USER | |
| TRUSTWORTHY | |

WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

| | |
|------------------------|---|
| | ▼ |
| MOINDEX | |
| PHYSICAL_ONLY | |
| REPAIR_ALLOW_DATA_LOSS | |
| REPAIR_FAST | |

WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

| | |
|--------------|---|
| | ▼ |
| MULTI_USER; | |
| ONLINE; | |
| OPEN; | |
| TRUSTWORTHY; | |

GO

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SINGLE_USER

The specified database must be in single-user mode to use one of the following repair options. Box 2: REPAIR_ALLOW_DATA_LOSS

REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.

Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql>

NEW QUESTION 281

- (Exam Topic 5)

You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.

You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

- A. Upgrade workspace1 to the Premium pricing tier.
- B. Configure a global init script for workspace1.
- C. Create a pool in workspace1.
- D. Create a cluster policy in workspace1.

Answer: C

Explanation:

You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.
 Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.
 Reference:
<https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html>

NEW QUESTION 286

- (Exam Topic 5)
 You have 50 Azure SQL databases.
 You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.
 What should you do?

- A. Create a diagnostic setting for the activity log that has the Security log enabled.
- B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
- C. Create an alert rule that uses a Metric signal type.
- D. Create an alert rule that uses an Activity Log signal type.

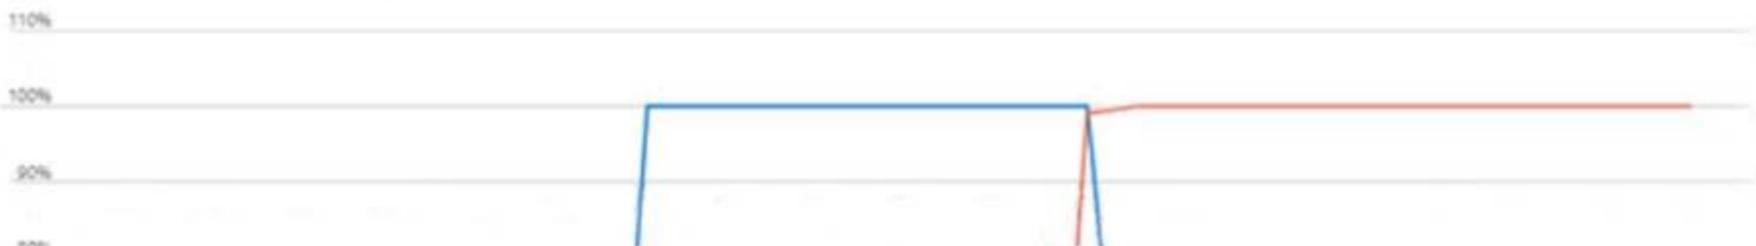
Answer: D

Explanation:

Activity log events - An alert can trigger on every event, or, only when a certain number of events occur. Reference:
<https://docs.microsoft.com/en-us/azure/azure-sql/database/alerts-insights-configure-portal>

NEW QUESTION 287

- (Exam Topic 5)
 You have an Azure SQL database named that contains a table named Table1. You run a query to bad data into Table1.
 The performance Of Table1 during the load operation are shown in exhibit.



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

- scale the resource
- use an elastic pool
- 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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

- scale the resource
- use an elastic pool
- 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

NEW QUESTION 292

- (Exam Topic 5)

You have an Azure SQL database named DBI that contains a nonclustered index named index1. End users report slow queries when they use index1. You need to identify the operations that are being performed on the index. Which dynamic management view should you use?

- A. sys.dm_exec_query_plan_stats
- B. Sys.dm_db_index_physical_stats
- C. Sys.dm_db_index_operational_stats
- D. Sys.dm_db_index_usage_stats

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 293

- (Exam Topic 5)

You have an Azure SQL managed instance that hosts multiple databases.

You need to configure alerts for each database based on the diagnostics telemetry of the database. What should you use?

- A. Azure SQL Analytics alerts based on metrics
- B. SQL Health Check alerts based on diagnostics logs
- C. SQL Health Check alerts based on metrics
- D. Azure SQL Analytics alerts based on diagnostics logs

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo>

NEW QUESTION 297

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource_semaphore waits.

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

NOTE:Each correct selection is worth one point.

```

SELECT
    [is_user_process]
    [wait_time]
    [wait_type]
SUM(wait_time) AS total_wait_time_ms
FROM sys.[dm_exec_query_stats]
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs>

NEW QUESTION 299

- (Exam Topic 5)

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1. You need to encrypt the data in Col1.

Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Create a database master key.
- Create a column master key.
- Open the symmetric key.
- Create a certificate.
- Update Col1.
- Create a symmetric key.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

Reference:

<https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/>

NEW QUESTION 304

- (Exam Topic 5)

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database. Which database option should set?

- A. PARAMETERIZATIONtoFORCED
- B. PARAMETERIZATIONtoSIMPLE
- C. Delayed Durability toForced
- D. READ_COMMITTED_SNAPSHOTtoON

Answer: D

Explanation:

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

- > The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.
- > The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql>

NEW QUESTION 309

- (Exam Topic 5)

You have a 50-TB Microsoft SQL Server database named DB1.

You need to reduce the time it takes to perform database consistency checks of DB1.

Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser>

NEW QUESTION 313

- (Exam Topic 5)

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:

- > 24 cores
- > 500 GB of storage
- > 124 GB of memory
- > 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

Answer: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcore-single-databases#business-critic>

NEW QUESTION 317

- (Exam Topic 5)

You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:

- > Support failovers that do not require client applications to change their connection strings.
- > Replicate the database to a secondary Azure region.
- > Support failover to the secondary region. What should you include in the recommendation?

- A. failover groups
- B. transactional replication
- C. Availability Zones
- D. geo-replication

Answer: A

Explanation:

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>

NEW QUESTION 319

- (Exam Topic 4)

You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data:

| | |
|-------------|---|
| | ▼ |
| Hash | |
| Replicated | |
| Round-robin | |

Table type to store promotional data:

| | |
|-------------|---|
| | ▼ |
| Hash | |
| Replicated | |
| Round-robin | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: Hash Scenario:

Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.

A hash distributed table can deliver the highest query performance for joins and aggregations on large tables. Box 2: Round-robin

Scenario:

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:

- > When you cannot identify a single key to distribute your data.
- > If your data doesn't frequently join with data from other tables.
- > When there are no obvious keys to join.

Reference:

<https://rajaniishkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/>

NEW QUESTION 321

- (Exam Topic 4)

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements.

What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table the has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

Answer: B

Explanation:

Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesidentity>

NEW QUESTION 322

- (Exam Topic 3)

Which counter should you monitor for real-time processing to meet the technical requirements?

- A. SU% Utilization
- B. CPU% utilization
- C. Concurrent users
- D. Data Conversion Errors

Answer: B

Explanation:

Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitor-tune-overview>

NEW QUESTION 326

- (Exam Topic 3)

Which windowing function should you use to perform the streaming aggregation of the sales data?

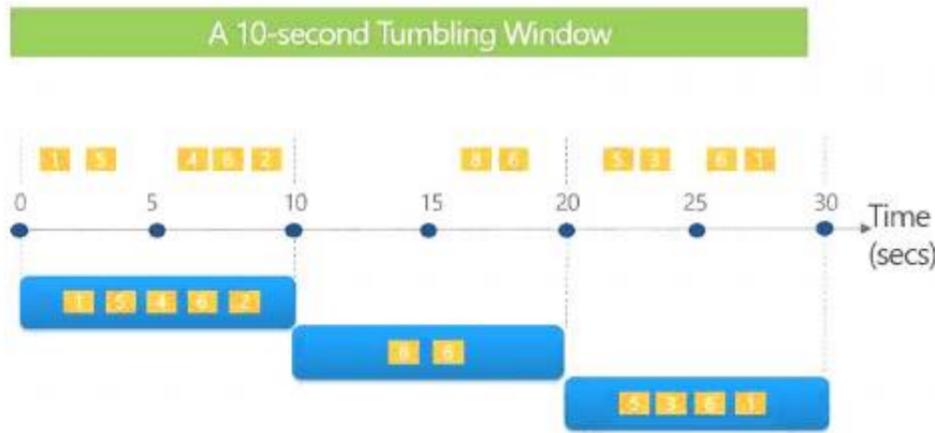
- A. Sliding
- B. Hopping
- C. Session
- D. Tumbling

Answer: D

Explanation:

Scenario: The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping. Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window. Timeline Description automatically generated

Tell me the count of Tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Reference:
<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun>

NEW QUESTION 329

- (Exam Topic 2)
 You are evaluating the role assignments.
 For each of the following statements, select Yes if the statement is true. Otherwise, select No.
 NOTE:Each correct selection is worth one point.

| Statements | Yes | No |
|---|-----------------------|-----------------------|
| DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio. | <input type="radio"/> | <input type="radio"/> |
| DBAGroup1 will be able to assign the SQL DB Contributor role to other users. | <input type="radio"/> | <input type="radio"/> |
| DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
 DBAGroup1 is member of the Contributor role.
 The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.
 Box 2: No
 Box 3: Yes
 DBAGroup2 is member of the SQL DB Contributor role.
 The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.

Reference:
<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

NEW QUESTION 333

- (Exam Topic 1)

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.

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

NOTE:Each correct selection is worth one point.

Purchasing model: ▼

| |
|--|
| Azure virtual machine reserved instances |
| DTU |
| vCore |

Deployment option: ▼

| |
|---|
| An Azure SQL Database elastic pool |
| An Azure SQL Database managed instance |
| A SQL Server Always On availability group |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DTU

Scenario:

- > The 30 new databases must scale automatically.
- > Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview> <https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview>

NEW QUESTION 336

- (Exam Topic 1)

You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

What should you include in the migration plan? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Azure Database Migration Service pricing tier: ▼

| |
|------------------|
| Standard 2-vCore |
| Standard 4-vCore |
| Premium 4-vCore |

Required Azure resource: ▼

| |
|--|
| A virtual network that has service endpoints |
| A VPN gateway |
| An Azure Logic app |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure Database Migration service Box 1: Premium 4-VCore

Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.

- > Minimize downtime during the migration of the SERVER1 databases.

Premium 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.

Reference: <https://azure.microsoft.com/pricing/details/database-migration/>

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online>

NEW QUESTION 340

- (Exam Topic 1)

You create all of the tables and views for ResearchDB1.

You need to implement security for ResearchDB1. The solution must meet the security and compliance requirements.

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

| Actions | Answer Area |
|--|-------------|
| Run the Always Encrypted wizard. | |
| Create an Azure Key Vault instance and generate a secret. | |
| Create an Azure Key Vault instance and configure an access policy. | ⬅️ ⬆️ |
| Create an Azure AD managed identity. | ➡️ ⬇️ |
| Register ResearchApp1 to Azure AD. | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=az>

NEW QUESTION 345

- (Exam Topic 1)

You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals. What should you include in the recommendation?

- A. For each customer, grant the customer ddl_admin to the existing schema.
- B. For each customer, create an additional schema and grant the customer ddl_admin to the new schema.
- C. For each customer, create an additional schema and grant the customer db_writerto the new schema.
- D. For each customer, grant the customer db_writerto the existing schema.

Answer: B

NEW QUESTION 346

- (Exam Topic 1)

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements. How should you collect and stream metrics? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Collect metrics from:

- The database only
- The elastic pool and the database
- The elastic pool only
- The server, the elastic pool, and the database

Stream metrics to:

- Azure Event Hubs
- Azure Log Analytics
- Azure Storage

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: The server, the elastic pool, and the database

Senario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform. Event hubs are able to handle custom metrics.

NEW QUESTION 349

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DP-300 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DP-300 Product From:

<https://www.2passeasy.com/dumps/DP-300/>

Money Back Guarantee

DP-300 Practice Exam Features:

- * DP-300 Questions and Answers Updated Frequently
- * DP-300 Practice Questions Verified by Expert Senior Certified Staff
- * DP-300 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DP-300 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year