

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 Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

- > Automatically deletes the logs at the end of each retention period
- > Minimizes storage costs

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

NOTE: Each correct selection is worth one point.

To minimize storage costs:

Store the infrastructure logs and the application logs in the Archive access tier.

Store the infrastructure logs and the application logs in the Cool access tier.

Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

Azure Data Factory pipelines

Azure Blob storage lifecycle management rules

Immutable Azure Blob storage time-based retention policies

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A picture containing text Description automatically generated

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier 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: Azure Blob storage lifecycle management rules

Blob storage lifecycle management offers a rich, rule-based policy that you can use to transition your data to the best access tier and to expire data at the end of its lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 4

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.

You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.

You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.

What should you run for the backup and where should you store the backup? 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:

Graphical user interface, application, Word Description automatically generated

Reference:

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

NEW QUESTION 5

- (Exam Topic 5)

You have an Azure subscription.

You need to deploy an Azure SQL managed instance by using an Azure Resource Manager (ARM) template. The solution must meet the following requirements:

The SQL managed instance must be assigned a unique identity.

The SQL managed instance must be available in the event of an Azure datacenter outage.

How should you complete the template? 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

```
{
  "type": "Microsoft.Sql/managedInstances",
  -
},
"identity": {
  "type": 
},
"dependsOn": [
  "[parameters('virtualNetworkName')]"
],
"properties": {
  "administratorLogin": "[parameters('administratorLogin')]",
  "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  parameters('virtualNetworkName'), parameters('subnetName'))]", "storageSizeInGB": 8192,
  "vCores": 80, "licenseType": "BasePrice",
   : "True"
}
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
{
  "type": "Microsoft.Sql/managedInstances",
  -
},
"identity": {
  "type": 
},
"dependsOn": [
  "[parameters('virtualNetworkName')]"
],
"properties": {
  "administratorLogin": "[parameters('administratorLogin')]",
  "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
  parameters('virtualNetworkName'), parameters('subnetName'))]", "storageSizeInGB": 8192,
  "vCores": 80, "licenseType": "BasePrice",
   : "True"
}
}
```

NEW QUESTION 6

- (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 reduce the use of table variables and temporary tables. 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 7

- (Exam Topic 5)

Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks. You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

Actions

Answer Area

- Deploy Grafana to an Azure virtual machine.
- Build a spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file.
- Create Dropwizard counters in the application code.
- Create a data source in Azure Monitor.
- Configure the Databricks cluster to use the Databricks monitoring library.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated with medium confidence
 Send application metrics using Dropwizard.
 Spark uses a configurable metrics system based on the Dropwizard Metrics Library.
 To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps: Step 1: Configure your Azure Databricks cluster to use the Databricks monitoring library.
 Prerequisite: Configure your Azure Databricks cluster to use the monitoring library. Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file
 Step 3: Create Dropwizard counters in your application code Create Dropwizard gauges or counters in your application code

NEW QUESTION 8

- (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 9

- (Exam Topic 5)

You have an Azure Data Factory that contains 10 pipelines.
 You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.
 What should you add to each pipeline?

- A. an annotation
- B. a resource tag

- C. a run group ID
- D. a user property
- E. a correlation ID

Answer: A

Explanation:

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.

Reference:

<https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/>

NEW QUESTION 10

- (Exam Topic 5)

You have an Azure SQL database named DB 1 in the General Purpose service tier. You need to monitor DB 1 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.

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:

Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql>

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 12

- (Exam Topic 5)

You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.

You need to create a surrogate key for the table. The solution must provide the fastest query performance. What should you use for the surrogate key?

- A. an IDENTITY column
- B. a GUID column
- C. a sequence object

Answer: A

Explanation:

Dedicated SQL pool supports many, but not all, of the table features offered by other databases. Surrogate keys are not supported. Implement it with an Identity column.

Reference:

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

NEW QUESTION 14

- (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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 17

- (Exam Topic 5)

You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:

- > ProductID
- > ItemPrice
- > LineTotal
- > Quantity
- > StoreID
- > Minute
- > Month
- > Hour
- > Year
- > Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.

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

NOTE: Each correct selection is worth one point.

df.write

.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour", "StoreID")

.mode ("append")

.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: .partitionBy Example:

df.write.partitionBy("y","m","d") mode(SaveMode.Append) parquet("/data/hive/warehouse/db_name.db/" + tableName) Box 2: ("Year","Month","Day","Hour","StoreID")

Box 3: .parquet("/Purchases") Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partiti>

NEW QUESTION 20

- (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 25

- (Exam Topic 5)

You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.

Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure. You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.

What should you include in the recommendation?

- A. Azure Backup
- B. Azure Information Protection
- C. Windows Server Failover Cluster
- D. Azure Site Recovery

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure>

NEW QUESTION 27

- (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 29

- (Exam Topic 5)

You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.

You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders.

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.

Create:

- A materialized view in DB1
- A security policy in the Orders table**
- Database scoped credentials in DB1

Control access to the rows by using:

- A masking rule**
- A table-valued function
- The CONTAINS predicate

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Create:

- A materialized view in DB1
- A security policy in the Orders table**
- Database scoped credentials in DB1

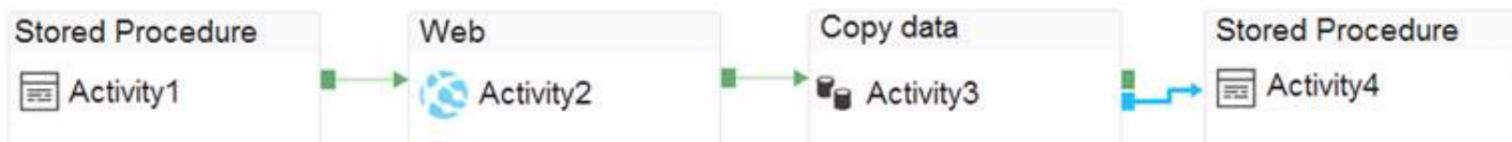
Control access to the rows by using:

- A masking rule**
- A table-valued function**
- The CONTAINS predicate

NEW QUESTION 31

- (Exam Topic 5)

You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

- > Operator: Greater than
- > Aggregation type: Total
- > Threshold value: 2
- > Aggregation granularity (Period): 5 minutes
- > Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

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
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: No

Just one failure within the 5-minute interval. Box 2: No

Just two failures within the 5-minute interval. Box 3: No

Just two failures within the 5-minute interval. Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

NEW QUESTION 34

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1.

Which isolation level should you configure?

- A. SERIALIZABLE
- B. SNAPSHOT
- C. READ UNCOMMITTED
- D. READ COMMITTED

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database>

NEW QUESTION 35

- (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 42

- (Exam Topic 5)

You have an Azure SQL Database managed instance. The instance starts experiencing performance issues.

You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.

What should you use?

- A. the Azure portal
- B. Extended Events
- C. Query Store
- D. dynamic management views

Answer: D

Explanation:

Reference:

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

NEW QUESTION 45

- (Exam Topic 5)

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

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

The database fails on Wednesday at 10:00.

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

Actions

Answer Area

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Answer Area

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

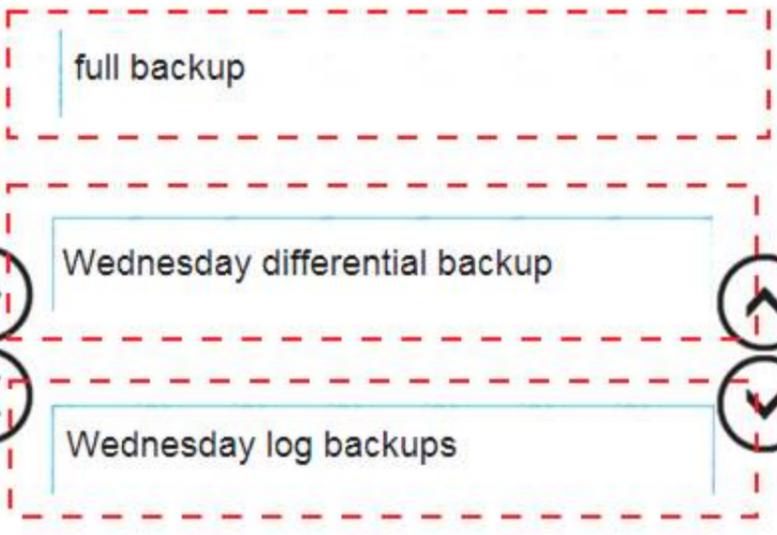
full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup



NEW QUESTION 48

- (Exam Topic 5)

You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

```
{
  "id":123,
  "address_housenumber": "19c",
  "address_line1": "Memory Lane",
  "applicant1_name": "Jane",
  "applicant2_name": "Dev"
}
```

The records contain two applicants at most.

You need to build a table that includes only the address fields.

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.

```

CREATE EXTERNAL TABLE
CREATE TABLE
CREATE VIEW
WITH (
    LOCATION = 'applications/',
    DATA_SOURCE = applications_ds,
    FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
    (BULK 'https://contoso1.dfs.core.windows.net/
    applications/year=*/*.parquet',
    FORMAT = 'PARQUET') AS [r]
GO

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: CREATE EXTERNAL TABLE

An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.

Syntax:

```
CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name }
( <column_definition> [ ,...n ] ) WITH (
```

```
LOCATION = 'folder_or_filepath', DATA_SOURCE = external_data_source_name, FILE_FORMAT = external_file_format_name
```

Box 2. OPENROWSET
 When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.

Example: AS

```
SELECT decennialTime, stateName, SUM(population) AS population FROM
OPENROWSET(BULK
'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/
FORMAT='PARQUET') AS [r]
GROUP BY decennialTime, stateName GO
```

Reference:

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

NEW QUESTION 53

- (Exam Topic 5)

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container. Which resource provider should you enable?

- A. Microsoft.EventHub
- B. Microsoft.EventGrid
- C. Microsoft.Sql
- D. Microsoft.Automation

Answer: B

Explanation:

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

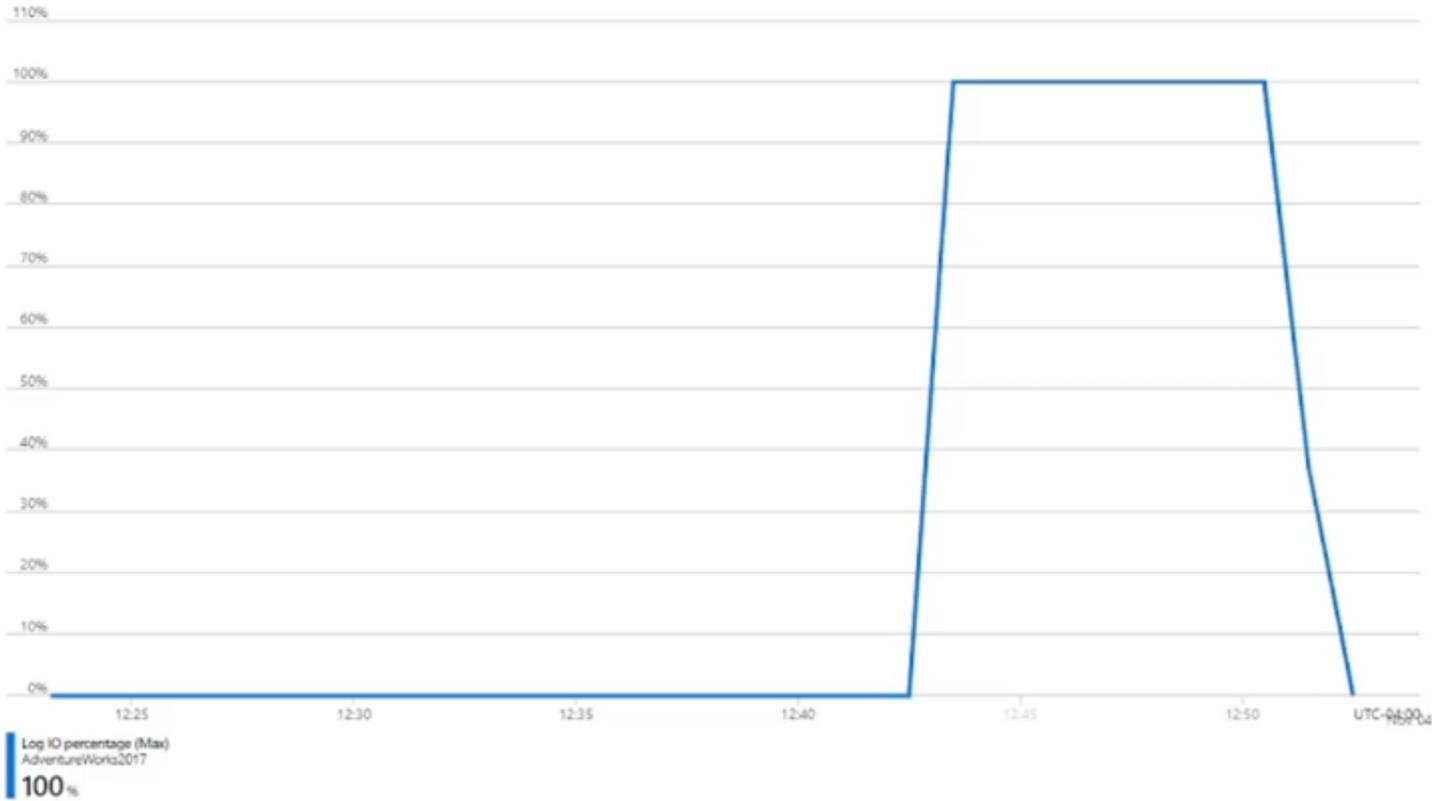
Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

NEW QUESTION 58

- (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 IO percentage. The solution must minimize costs. What should you do?

- A. Increase the number of vCores.
- B. Change Recoverymodel to Simple.
- C. Perform a checkpoint operation.
- D. Change Service tier to Business Critical.

Answer: D

NEW QUESTION 61

- (Exam Topic 5)

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

You need to receive voice message notifications when a maintenance event affects any of the 10 regions. The solution must minimize administrative effort. What should you do?

- A. From the Azure portal, create a service health alert.
- B. From the Azure portal, create an Azure Advisor operational excellence alert.
- C. From Microsoft SQL Server Management Studio (SSMS), configure a SQL Server agent job.
- D. From the Azure portal, configure an activity log alert.

Answer: C

NEW QUESTION 65

- (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 69

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled.

During a database audit, you discover that none of the fields contain more than 150 characters. You need to ensure that you can apply page compression to Table1.

What should you do?

- A. Configure the columns as sparse.
- B. Change the column type to nvarchar(MAX).
- C. Change the column type to varchar(MAX).
- D. Change the column type to varchar(200).

Answer: D

Explanation:

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/> <https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/>

NEW QUESTION 72

- (Exam Topic 5)

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest. What should you enable?

- A. Transparent Data Encryption (TDE)
- B. Advanced Data Security for this database
- C. Always Encrypted for all columns
- D. Secure transfer required

Answer: A

Explanation:

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview>

NEW QUESTION 73

- (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

▼
DATABASE
FILE
LOG

 DB1 FROM

▼
DISK = N'\\NAS01\SSQLBackups\DB1.bak';
TAPE = N'\\Tape0'
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 76

- (Exam Topic 5)

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

You need to identify which database queries consume the most resources. Which tool should you use?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

Answer: A

Explanation:

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

NEW QUESTION 78

- (Exam Topic 5)

You have an Azure subscription that contains an Azure SQL managed instance named SQLMi1 and a SQL Agent job named Backupdb. Backupdb performs a daily backup of the databases hosted on SQLMi1.

You need to be notified by email if the 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.

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 SQL Server Agent alert.
- Create an operator.
- Create an extended event.
- Enable Database Mail.
- Add a failure notification to the job.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance>

NEW QUESTION 81

- (Exam Topic 5)

You have two instances of SQL Server on Azure virtual Machines named VM1 and VM2. VM1 hosts a database named db1.

You plan to create a database availability group (DAG) for db1. The solution must use certificate authentication between VM1 and VM2.

You need to configure authentication for the outbound connections of VM1.

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

- From the master database, run CREATE_ENDPOINT.
- From db1, run CREATE_CERTIFICATE.
- From the master database, run CREATE_MASTER_KEY_ENCRYPTION_BY_PASSWORD.
- From the master database, run CREATE_CERTIFICATE.
- From db1, run CREATE_ENDPOINT.

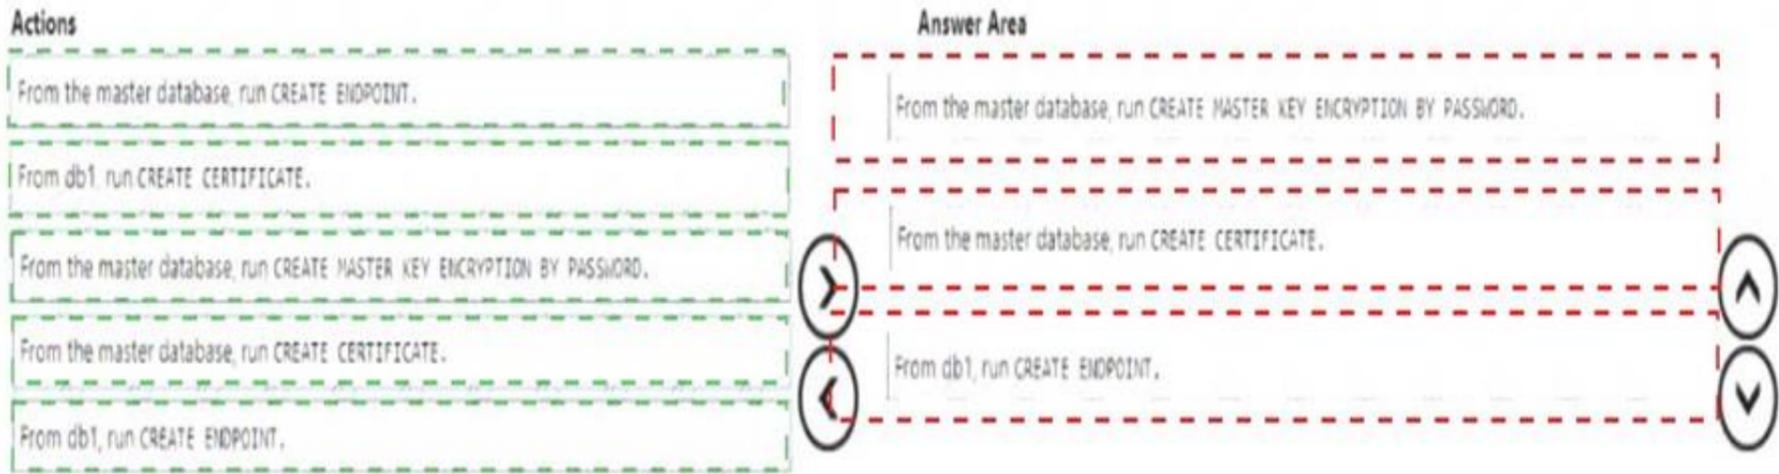
Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 86

- (Exam Topic 5)

You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases. You need to update the column and index statistics for the databases.

What should you use?

- A. an Azure Automation runbook
- B. a SQL Agent job
- C. Azure SQL Analytics
- D. automatic tuning in Azure SQL Database

Answer: A

Explanation:

Reference:

<https://www.sqlshack.com/automate-azure-sql-database-indexes-and-statistics-maintenance/>

NEW QUESTION 88

- (Exam Topic 5)

You have a single availability set that contains two SQL Server on Azure Virtual Machines instances.

The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1.

You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:

- Provide an availability SLA
- Minimize costs.

What should you create?

- A. a virtual network name (VNN) resource
- B. a Basic Azure Load Balancer
- C. a distributed network name (DNN) resource
- D. an Azure Standard Load Balancer

Answer: C

NEW QUESTION 90

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines.

You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The solution must use the principle of least privilege.

Which role should you assign to User1?

- A. sysadmin
- B. SQLAgentUserRole
- C. SQLAgentReaderRole
- D. SQLAgentOperatorRole

Answer: A

NEW QUESTION 92

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine that contains a database named Db1. You need to enable automatic tuning for Db1.

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

NOTE:Each correct selection is worth one point.

ALTER DATABASE [Db1]

```

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
    
```

GO

ALTER DATABASE [Db1]

```

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)
SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)
SET AUTOMATIC_TUNING=AUTO
SET QUERY_STORE=OFF
SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)
SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)
    
```

GO

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SET AUTOMATIC_TUNING = AUTO

To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option. Setting it to OFF will also override any setting that database inherited and disable the tuning option.

Reference:

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

NEW QUESTION 95

- (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 DBI. 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 99

- (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: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

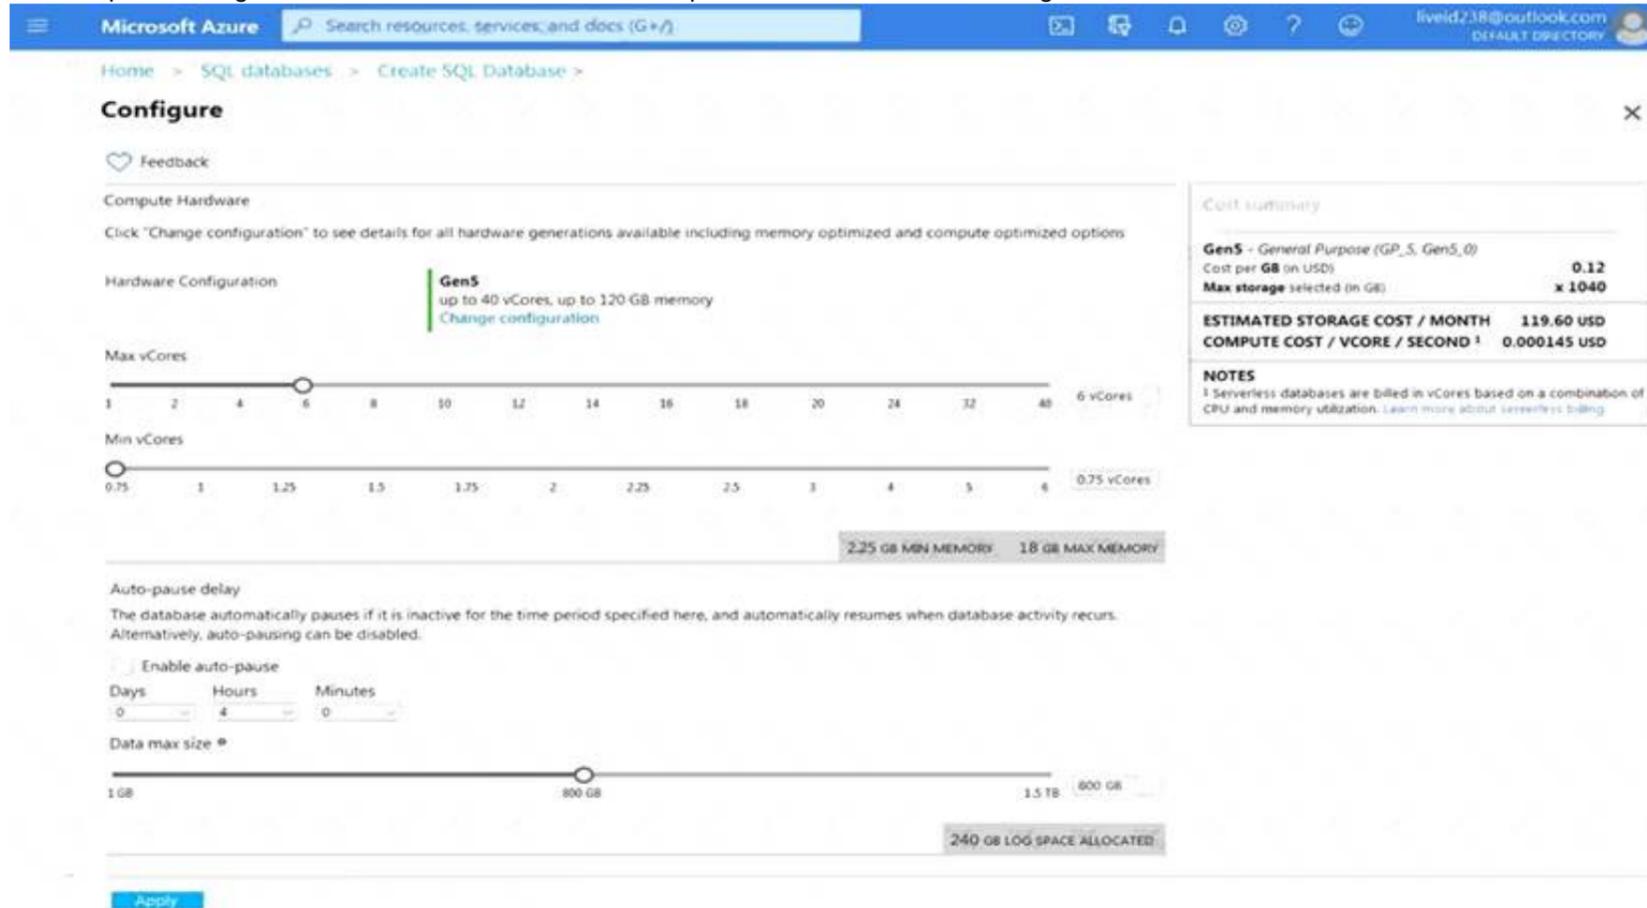
Reference:

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

NEW QUESTION 102

- (Exam Topic 5)

You are provisioning an Azure SQL database in the Azure portal 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.

After four hours of inactivity, the database requires [answer choice] to resume operations for new activities.

no extra time
 up to 10 minutes
 up to one minute

The database configuration reduces the cost of [answer choice] usage patterns.

intermittent and unpredictable
 regular and high
 steady and low

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 107

- (Exam Topic 5)

You have an Azure subscription that contains a storage account named databasebackups. You have an Azure SQL managed instance named DB1.

You need to back up DB1 to databasebackups.

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

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE'
                'DatabaseBackups'
                'KeyVault1'
                'SHARED ACCESS SIGNATURE'
SECRET = 'spr&st=2023-02-02T19:23:08Z&se=2033-02-
02T19:30:08Z&spr=https&sv=2021-06-
08&sr=b&sig=B%2FxEYQi0C%4BqyYCeqlwHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL =
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
      CHECKSUM
      COMPRESSION
      COPY_ONLY
      DIFFERENTIAL
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE'
                'DatabaseBackups'
                'KeyVault1'
                'SHARED ACCESS SIGNATURE'
SECRET = 'spr&st=2023-02-02T19:23:08Z&se=2033-02-
02T19:30:08Z&spr=https&sv=2021-06-
08&sr=b&sig=B%2FxEYQi0C%4BqyYCeqlwHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL =
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
      CHECKSUM
      COMPRESSION
      COPY_ONLY
      DIFFERENTIAL
```

NEW QUESTION 112

- (Exam Topic 5)

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8. You need to deploy a failover cluster instance (FCI) to VM1 and VM2. The solution must eliminate the need for the following:

- A distributed network name (DNN)
 - A load balancer
- What should you do?

- A. Deploy VM1 and VM2 to a single proximity placement group.
- B. Deploy VM1 and VM2 to different proximity placement groups in the same Azure region.
- C. Connect VM1 and VM2 to a single subnet.
- D. Connect VM1 and VM2 to different subnets on a single virtual network.

Answer: D

NEW QUESTION 113

- (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 116

- (Exam Topic 5)

You have an Azure subscription that contains 50 instances of SQL Server on Azure Virtual Machines. The instances host 500 Azure SQL databases. You need to ensure that all the databases have the same configuration. The solution must meet the following requirements:

- Auditing must be enabled.
- Azure Defender must be enabled.
- Public network access must be disabled.
- Administrative effort must be minimized.

Which two resources should you create in the subscription? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. an Azure Policy assignment
- B. an Azure Automation account
- C. an Azure Policy initiative
- D. an Azure Automation runbook
- E. an Azure Policy definition

Answer: CE

NEW QUESTION 118

- (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 120

- (Exam Topic 5)

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

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
Change the Active Directory Admin on TestServer1	
Change the server name and related variables in the templates	
From the database project, deploy the database schema and permissions	⤴
Add IP addresses to the firewall	⤵
From the Azure portal, export the Azure Resource Manager templates	⤴
From the Azure portal, deploy the templates.	⤵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Change the Active Directory Admin on TestServer1	
Change the server name and related variables in the templates	From the Azure portal, export the Azure Resource Manager templates
From the database project, deploy the database schema and permissions	Change the server name and related variables in the templates
Add IP addresses to the firewall	From the Azure portal, deploy the templates.
From the Azure portal, export the Azure Resource Manager templates	
From the Azure portal, deploy the templates.	From the database project, deploy the database schema and permissions

NEW QUESTION 122

- (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 127

- (Exam Topic 5)

You have the following Azure Data Factory pipelines:

- Ingest Data from System1
- Ingest Data from System2

- > Populate Dimensions
- > Populate Facts

Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2. Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours. What should you do to schedule the pipelines for execution?

- A. Add a schedule trigger to all four pipelines.
- B. Add an event trigger to all four pipelines.
- C. Create a parent pipeline that contains the four pipelines and use an event trigger.
- D. Create a parent pipeline that contains the four pipelines and use a schedule trigger.

Answer: D

Explanation:

Reference:
<https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/>

NEW QUESTION 132

- (Exam Topic 5)

You have SQL Server 2019 on 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 provide the lowest latency for tempdb. What is the total number of data files that tempdb should contain?

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

Answer: C

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference:
<https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

NEW QUESTION 137

- (Exam Topic 5)

You have an Azure SQL Database instance named DatabaseA on a server named Server1. You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication. You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials. 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
On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;		
On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';	⬅	⬆
On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];	➡	⬆
On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';		
On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
 Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:
 CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
 You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE

LOGIN command.

Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

CREATE USER readonlyuser FROM LOGIN readonlylogin;

Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]

Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the db_datareader role.

EXEC sp_addrolemember 'db_datareader', 'readonlyuser'; Reference:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

NEW QUESTION 141

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool.

You run PDW_SHOWSPACEUSED('dbo.FactInternetSales'); and get the results shown in the following table.

ROWS	RESERVED_SPACE	DATA_SPACE	INDEX_SPACE	UNUSED_SPACE	PDW_NODE_ID	DISTRIBUTION_ID
694	2776	616	48	2112	1	1
407	2704	576	48	2080	1	2
53	2376	512	16	1848	1	3
58	2376	512	16	1848	1	4
168	2632	528	32	2072	1	5
195	2696	536	32	2128	1	6
5995	3464	1424	32	2008	1	7
0	2232	496	0	1736	1	8
264	2576	544	40	1992	1	9
3008	3016	960	32	2024	1	10
...
1550	2832	752	48	2032	1	50
1238	2832	696	40	2096	1	51
192	2632	528	32	2072	1	52
1127	2768	680	48	2040	1	53
1244	3032	704	64	2264	1	54
409	2632	568	32	2032	1	55
0	2232	496	0	1736	1	56
1437	2832	728	40	2064	1	57
0	2232	496	0	1736	1	58
384	2632	560	32	2040	1	59
225	2768	544	40	2184	1	60

Which statement accurately describes the dbo.FactInternetSales table?

- A. The table contains less than 10,000 rows.
- B. All distributions contain data.
- C. The table uses round-robin distribution
- D. The table is skewed.

Answer: D

Explanation:

The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.

Note: SHOWSPACEUSED displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database.

This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.

ROUND_ROBIN distributed tables should not be skewed. Data is distributed evenly across the nodes by design.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu> <https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distrib>

NEW QUESTION 143

- (Exam Topic 5)

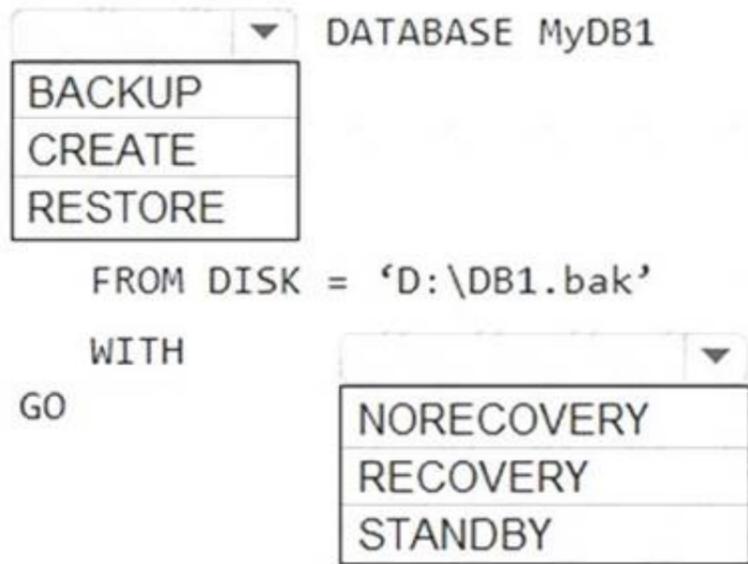
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.

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 144

- (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 149

- (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 154

- (Exam Topic 5)

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales. You need to initiate a backup of Sales. 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.

BACKUP DATABASE Sales

TO DISK = \\BackupSystem\BackupDisk1\Sales.bak'
TO DISK = 'X:\BAK\Sales.bak'
TO 'Sales_Backup'
TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

WITH COPY_ONLY;
WITH ENCRYPTION;
WITH FILE_SNAPSHOT;
WITH NO_TRUNCATE

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak' Native database backup in Azure SQL Managed Instance. You can backup any database using standard BACKUP T-SQL command: BACKUP DATABASE tpcc2501 TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak' WITH COPY_ONLY
 Box 2: WITH COPY_ONLY
 Reference:
<https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-insta>

NEW QUESTION 159

- (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 162

- (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 163

- (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
  "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 166

- (Exam Topic 5)

You have SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

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
Export a built-in policy.	
Create a custom policy based on a condition.	
Create a custom condition based on a built-in facet.	⬅️ ⬆️
View the policy history.	➡️ ⬇️
Import a policy file.	
Run a policy evaluation.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated
 Reference:
<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-bas>

NEW QUESTION 168

- (Exam Topic 5)

You have an Azure SQL database that contains a table named factSales. FactSales contains the columns shown in the following table.

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.
 Which type of compression provides the greatest space reduction for the database?

- A. page compression
- B. row compression

- C. columnstore compression
- D. columnstore archival compression

Answer: D

Explanation:

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression>

NEW QUESTION 171

- (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 173

- (Exam Topic 5)

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

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

You need to configure a connection between VM1 and MI1. The solution must meet the following requirements:

- The connection must be encrypted.
- Network latency must be minimized. What should you implement?

- A. virtual network peering
- B. private endpoints
- C. service endpoints
- D. a site-to-site VPN

Answer: B

NEW QUESTION 175

- (Exam Topic 5)

You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

- A. Azure Event Hubs Standard
- B. Azure Stream Analytics
- C. Azure Data Factory
- D. Azure Event Hubs Dedicated

Answer: D

Explanation:

The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.

Reference:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

NEW QUESTION 176

- (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 run theRemove-AzSqlDatabasePowerShell cmdlet for Database1 on Server2. You run theRestore-AzSqlDatabasePowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

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

NEW QUESTION 178

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication. You need to migrate SQL1 to Azure. Which service should you use?

- A. Azure SQL Edge
- B. Azure SQL Database
- C. SQL Server on Azure Virtual Machines
- D. Azure SQL Managed instance

Answer: C

NEW QUESTION 181

- (Exam Topic 5)

You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features. You plan to migrate to Azure SQL. Which service should you use?

- A. Azure SQL Database
- B. SQL Server on an Azure Virtual Machine
- C. Azure SQL Managed Instance
- D. Azure Database for MySQL

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-overview>

NEW QUESTION 183

- (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 186

- (Exam Topic 5)

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1. What should you do?

- A. Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property.
- B. Use Transact-SQL to enable the force_last_good_plan option.
- C. From the Azure portal, configure automatic tuning.

Answer: B

NEW QUESTION 188

- (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 191

- (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 194

- (Exam Topic 5)

You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2 . SQLMI2 contains a database named DB1 and a user named User1.

User1 drops DB1.

You need to perform a point-in-time restore of DB1 to SQLMI2.

- A. Azure CLI
- B. Transact-SQL
- C. The Azure portal
- D. Azure PowerShell

Answer: C

NEW QUESTION 198

- (Exam Topic 5)

You have a SQL Server on Azure Virtual Machines instance named VM1 . You run the following query.

```
BACKUP LOG DB1 TO DISK = '\\File1\SQLBackups\DB1.trn'
WITH NORECOVERY, COPY_ONLY, CONTINUE_AFTER_ERROR;
GO
```

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

Answer Area

Statements	Yes	No
The log file will be truncated.	<input type="radio"/>	<input type="radio"/>
DB1 will be placed in an offline state.	<input type="radio"/>	<input type="radio"/>
You are performing a tail-log backup.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The log file will be truncated.	<input checked="" type="radio"/>	<input type="radio"/>
DB1 will be placed in an offline state.	<input type="radio"/>	<input checked="" type="radio"/>
You are performing a tail-log backup.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION 202

- (Exam Topic 5)

You configure backups 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 that 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]**

After the 52nd weekly backup runs, there will be **[answer choice]**

- 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.
- 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.

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

NEW QUESTION 206

- (Exam Topic 5)

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Source spid1595

Message
The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'
```

You need to ensure That db1 can process transactions.

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

- Remove db1 from the availability group.
- Shrink the transaction log file.
- Add db1 back to the availability group.

NEW QUESTION 210

- (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 211

- (Exam Topic 5)

You have an Azure SQL database.

You need to implement a disaster recovery solution that meets the following requirements:

- Minimizes how long it takes to recover the database if a datacenter fails
- Minimizes administrative effort

What should you include in the solution?

- A. Azure Backup
- B. active geo-replication
- C. Azure Site Recovery
- D. auto-failover groups

Answer: D

NEW QUESTION 214

- (Exam Topic 5)

Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection. You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database. What should you use?

- A. Azure Migrate
- B. Data Migration Assistant (DMA)
- C. Azure Data BOX
- D. Azure Database Migration Service

Answer: C

Explanation:

<https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Mig>

NEW QUESTION 218

- (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 220

- (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 223

- (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 228

- (Exam Topic 5)

You have the following Azure Resource Manager template.

```

...
  "variable": {
    "serverName": "azsqlserver0001"
  },
  "resources": [
    {
      "name": "[variables('serverName')]",
      "type": "Microsoft.Sql/servers",
      "apiVersion": "2019-06-01-preview",
      "location": "[parameters('location')]",
      "properties": {
        "administratorLogin": "[parameters('administratorLogin')]",
        "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
        "version": "12.0"
      },
    },
    {
      "name": "[concat(variables('serverName'),'/',parameters('databaseName'))]",
      "type": "Microsoft.Sql/servers/databases",
      "apiVersion": "2020-08-01-preview",
      "location": "[parameters('location')]",
      "kind": "v12.0"
      "sku": {
        "name": "Standard",
        "tier": "Standard",
        "capacity": 10
      },
      "dependsOn": [
        "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
      ],
      "properties": {
      },
      "resources": [
      ]
    }
  ]
}
]
}
...

```

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
The template deploys a serverless Azure SQL database.	<input type="radio"/>	<input type="radio"/>
The template deploys a database to an Azure SQL Database managed instance.	<input type="radio"/>	<input type="radio"/>
The pricing tier of the database deployment is based on DTUs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A screenshot of a computer Description automatically generated with low confidence

Reference:

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

NEW QUESTION 229

- (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 Appl. Appl 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.

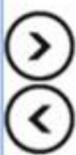
Actions

- 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:

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

Create a resource pool that has the following configurations.

```

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50
    
```

(Left arrow icon)

Create a workload group.

Create a classifier function in the master database.

NEW QUESTION 232

- (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 executes mapping data flow, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not a mapping flow, 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 233

- (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: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

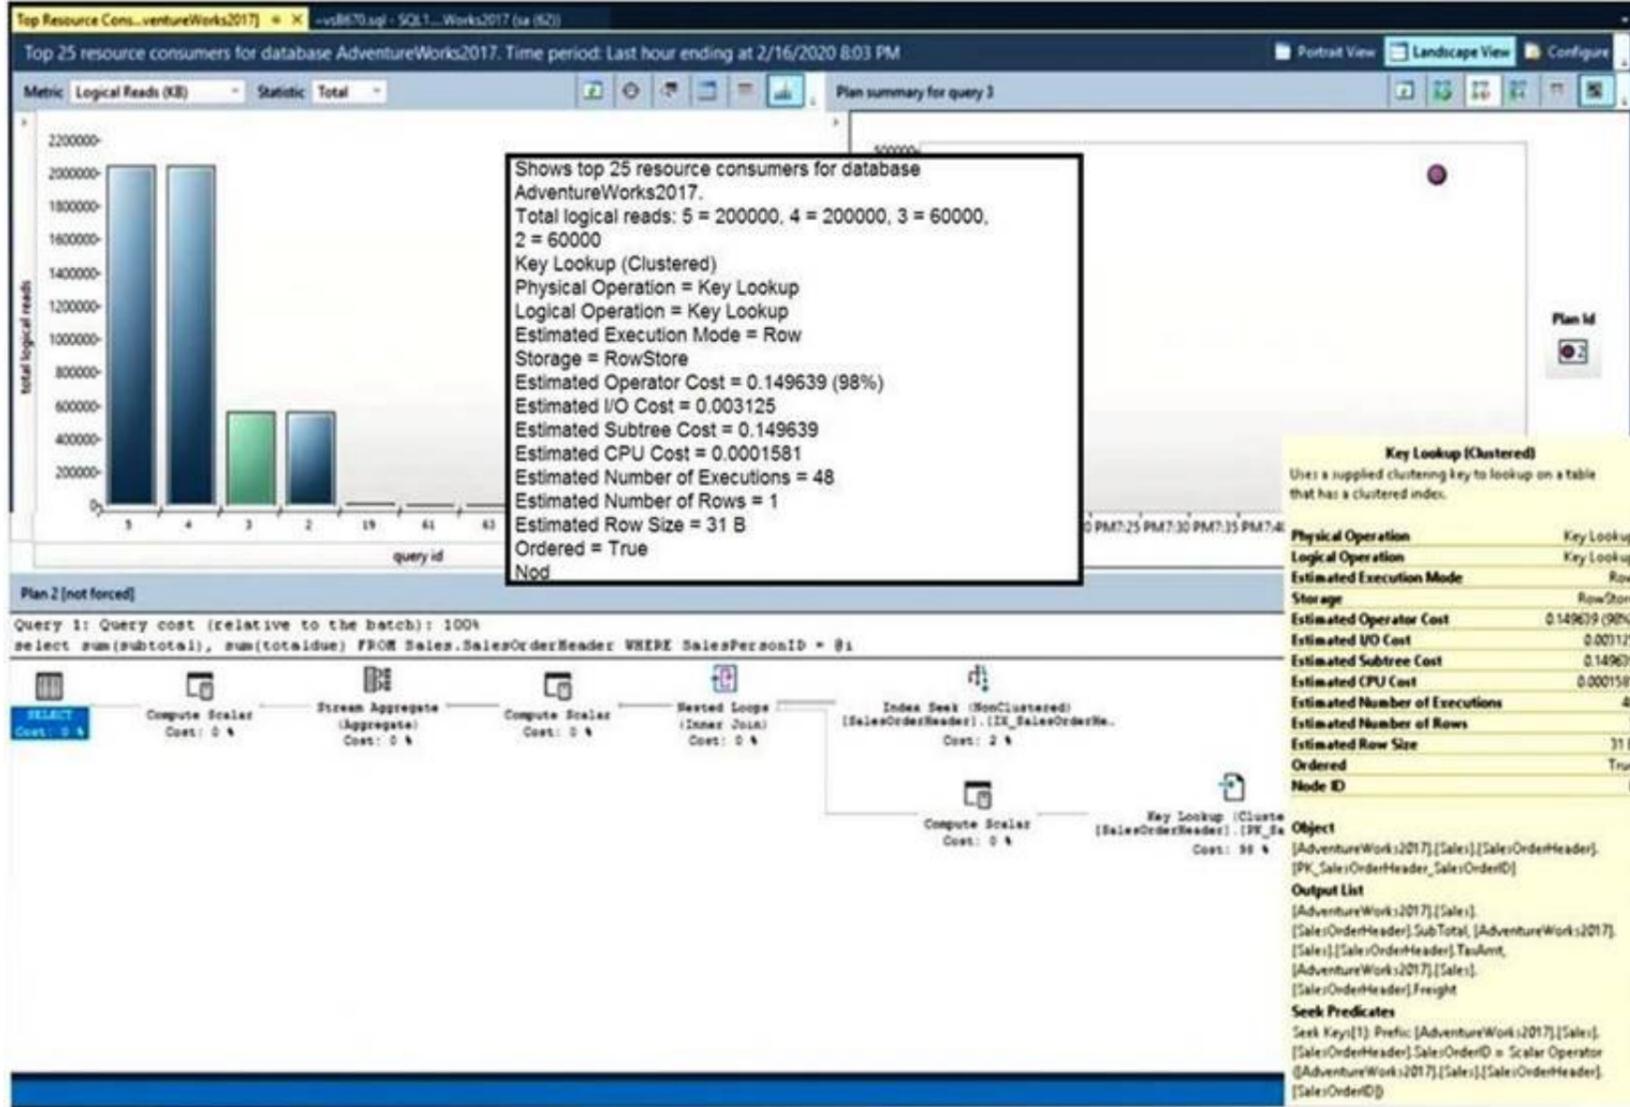
<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

NEW QUESTION 238

- (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 243

- (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

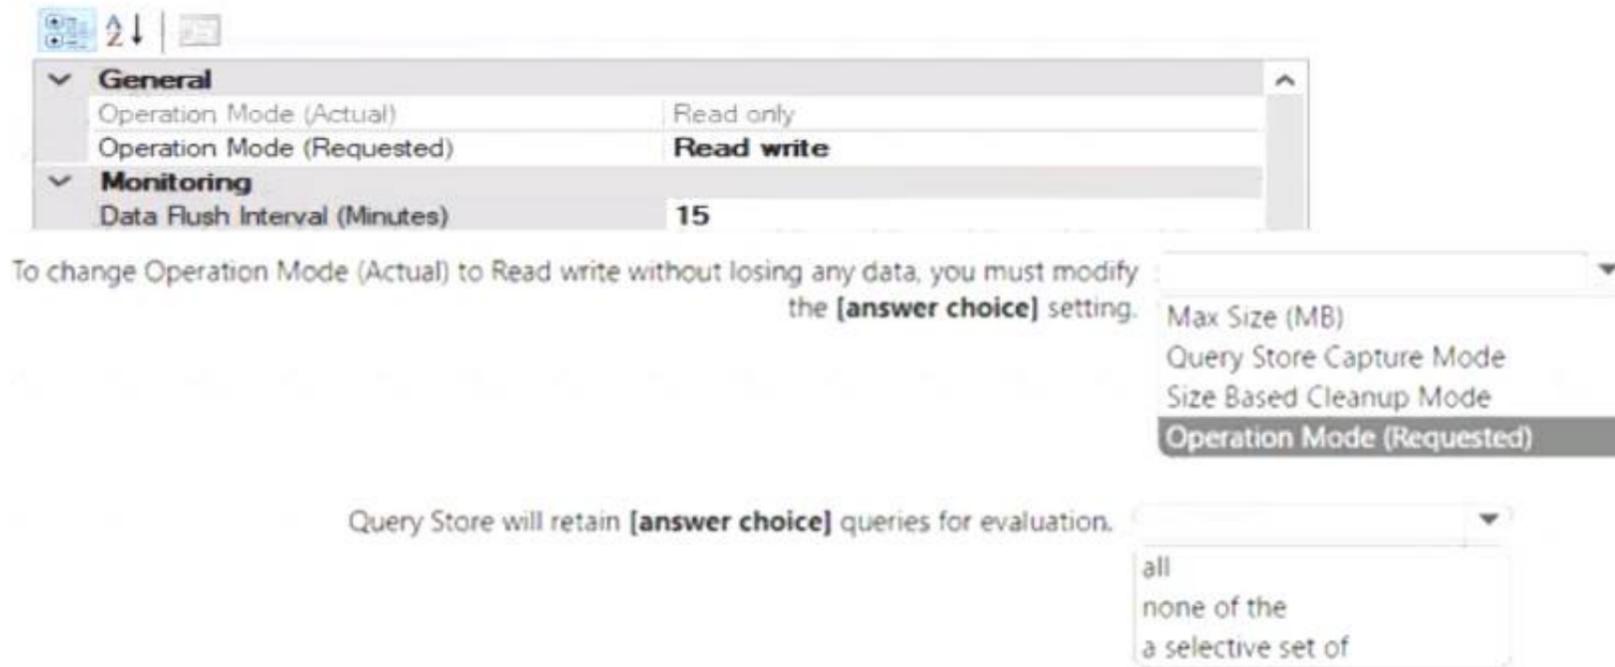
After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.

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

NEW QUESTION 245

- (Exam Topic 5)

You have a database on a SQL Server on Azure Virtual Machines instance. The current state of Query Store for the database is shown in the following exhibit.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text Description automatically generated

NEW QUESTION 249

- (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/azure-sql/database/secure-database-tutorial>

NEW QUESTION 250

- (Exam Topic 5)

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1. You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1. The on-premises network connects to VNET1 by using an ExpressRoute connection. You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service. You need to configure VNET1 to support the migration. What should you do?

- A. Configure service endpoints.
- B. Configure virtual network peering.
- C. Deploy an Azure firewall.
- D. Configure network security groups (NSGs).

Answer: A

Explanation:

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

NEW QUESTION 251

- (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 253

- (Exam Topic 5)

You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.

ADF1 contains the following pipelines:

- > P1: Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account
- > P2: Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2

You need to configure P1 and P2 to maximize parallelism and performance.

Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

P1:

P2:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

P1: Set the Partition option to Dynamic Range.

The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel. P2: Set the Copy method to PolyBase
 Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data

stores, including Azure Blob storage and Data Lake Store. (Polybase supports Azure Blob storage and Azure Data Lake Store by default.)

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse> <https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse>

NEW QUESTION 255

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table. 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.

CREATE EXTERNAL DATA SOURCE source1

WITH

(LOCATION = 'https://account1.  .core.windows.net',

blob
dfs
table

PUSHDOWN = ON
TYPE = BLOB_STORAGE
TYPE = HADOOP

)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, table Description automatically generated

Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi

WITH (LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/', TYPE = HADOOP)

Box 2: HADOOP

Reference:

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

NEW QUESTION 258

- (Exam Topic 5)

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

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events. You need to correlate the session events with events captured by Event Tracing for Windows (ETW). What should you do for session1?

- A. Modify the Set Session Event Filters settings.
- B. Add a target.
- C. Add an action.
- D. Modify the Specify Session Data Storage settings.

Answer: B

NEW QUESTION 262

- (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

- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.
- Create a contained database user.
- Create a resource pool.

Answer Area



- 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 264

- (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 General Purpose service tier and failover groups. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

Reference:

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

NEW QUESTION 265

- (Exam Topic 5)

You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.

You need to recommend a deployment solution that meets the following requirements:

- Provides a Service Level Agreement (SLA) of at least 99.95%
- Replicates databases in the same group synchronously
- Minimizes the latency of database writes What should you recommend?

- A. Create a proximity group and an availability se
- B. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group.
- C. Create two proximity groups and a single availability se
- D. Deploy both virtual machines to the availability se
- E. Add one virtual machine to each proximity group.
- F. Create two proximity groups and two availability set
- G. Deploy each virtual machine to a unique availability se
- H. Add one virtual machine to each proximity group.
- I. Create a proximity group and two availability set
- J. Deploy each virtual machine to a unique availability se
- K. Add both virtual machines to the proximity group.

Answer: A

NEW QUESTION 267

- (Exam Topic 5)

You have an Azure virtual machine based on a custom image named VM1. VM1 hosts an instance of Microsoft SQL Server 2019 Standard. You need to automate the maintenance of VM1 to meet the following requirements: Automate the patching of SQL Server and Windows Server. Automate full database backups and transaction log backups of the databases on VM1. Minimize administrative effort. What should you do first?

- A. Enable a system-assigned managed identity for VM1
- B. Register VM1 to the Microsoft.Sql resource provider
- C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1
- D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Answer: D

Explanation:

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaaSExtension) runs on Azure virtual machines to automate administration tasks. The SQL Server IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider. Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extensionauto>

NEW QUESTION 272

- (Exam Topic 5)

You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2. You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication. You need to ensure that App1 can access DB2. The solution must meet the following requirements:

- App1 must be able to view only DB2.
- Administrative effort must be minimized. What should you create?

- A. a contained database user for App1 on DB2
- B. a login for App1 on Server1
- C. a contained database user from an external provider for App1 on DB2
- D. a contained database user from a Windows login for App1 on DB2

Answer: D

Explanation:

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-databa>

NEW QUESTION 276

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- Clustered columnstore indexes
- Automatic tuning
- Change tracking
- PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

- A. Clustered columnstore indexes
- B. PolyBase
- C. Change tracking
- D. Automatic tuning

Answer: B

Explanation:

This table lists the key features for PolyBase and the products in which they're available.

Feature	SQL Server (Beginning with 2016)	Azure SQL Database	Azure Synapse Analytics	Parallel Data Warehouse
Query Hadoop data with Transact-SQL	Yes	No	No	Yes
Import data from Hadoop	Yes	No	No	Yes
Export data to Hadoop	Yes	No	No	Yes
Query, import from, export to Azure HDInsight	No	No	No	No
Push down query computations to Hadoop	Yes	No	No	Yes
Import data from Azure Blob storage	Yes	Yes*	Yes	Yes
Export data to Azure Blob storage	Yes	No	Yes	Yes
Import data from Azure Data Lake Store	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft BI tools	Yes	No	Yes	Yes

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary>

NEW QUESTION 277

- (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 279

- (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 in DB1. 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.dladmin` 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 280

- (Exam Topic 5)

You have an Azure subscription.

You plan to deploy an instance of SQL Server 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 283

- (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 286

- (Exam Topic 5)

You are building an Azure virtual machine. You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS. You need to configure storage for the virtual machine to support the database.

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

Actions	Answer Area
a virtual disk that uses the stripe layout	
a virtual disk that uses the mirror layout	
a volume	⤴
a virtual disk that uses the simple layout	⤵
a storage pool	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Follow these same steps to create striped virtual disk:

- > Create Log Storage Pool.
- > Create Virtual Disk
- > Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.

Box 3: a volume Reference:

<https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/>

NEW QUESTION 290

- (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 294

- (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 VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- > Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: B

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 298

- (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 303

- (Exam Topic 5)

You have an Azure SQL database named db1 on a server named server1. You use Query Performance Insight to monitor db1. You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible. Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE

Value:

1
60
CUSTOM
ON

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 307

- (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 only ReadWrite

NEW QUESTION 312

- (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

GO

	▼
OFFLINE	
ONLINE	
SINGLE_USER	
TRUSTWORTHY	

WITH ROLLBACK IMMEDIATE;

DBCC CHECKDB ('DB1',

GO

	▼
MOINDEX	
PHYSICAL_ONLY	
REPAIR_ALLOW_DATA_LOSS	
REPAIR_FAST	

WITH NO_INFOMSGS;

ALTER DATABASE [DB1] SET

GO

	▼
MULTI_USER;	
ONLINE;	
OPEN;	
TRUSTWORTHY;	

- 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 314

- (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 318

- (Exam Topic 5)

You are creating a managed data warehouse solution on Microsoft Azure.

You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.

You need to configure Azure Synapse Analytics to receive the data.

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 an external data source for Azure Blob storage.
- Create a master key on database.
- Enable Transparent Data Encryption.
- Create the external table FactSalesOrderDetails.
- Load the data to a staging table.
- Create an external file format to map the parquet files.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated
 To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.

Step 1: Create a master key on database.

* 1. Create a master key on the database. The master key is required to encrypt the credential secret. (Create a database scoped credential for Azure blob storage.)

Step 2: Create an external data source for Azure Blob storage.

* 2. Create an external data source with CREATE EXTERNAL DATA SOURCE.. Step 3: Create an external file format to map the parquet files.

* 3. Create an external file format with CREATE EXTERNAL FILE FORMAT. Step 4. Create an external table FactSalesOrderDetails

* 4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage>

NEW QUESTION 322

- (Exam Topic 5)

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

- A. Input Deserialization Errors
- B. Late Input Events
- C. Early Input Events
- D. Watermark delay

Answer: D

Explanation:

The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.

The watermark delay metric can rise due to:

* 1. Not enough processing resources in Stream Analytics to handle the volume of input events.

* 2. Not enough throughput within the input event brokers, so they are throttled.

* 3. Output sinks are not provisioned with enough capacity, so they are throttled. Reference:

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

NEW QUESTION 325

- (Exam Topic 5)

You have an Azure SQL database named that contains a table named Table1. You run a query to load 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 OLTP durable
- an In-Memory OLTP 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 OLTP durable
- an In-Memory OLTP non durable

NEW QUESTION 330

- (Exam Topic 5)

You have an Azure SQL database named DB1 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 331

- (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 336

- (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 
      
      
      SUM(wait_time) AS total_wait_time_ms
FROM sys.
      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 340

- (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 341

- (Exam Topic 5)

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

- A. An Azure SQL Database elastic pool
- B. SQL Server on Azure virtual machines
- C. an Azure SQL Database managed instance
- D. Azure SQL databases

Answer: A

Explanation:

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.

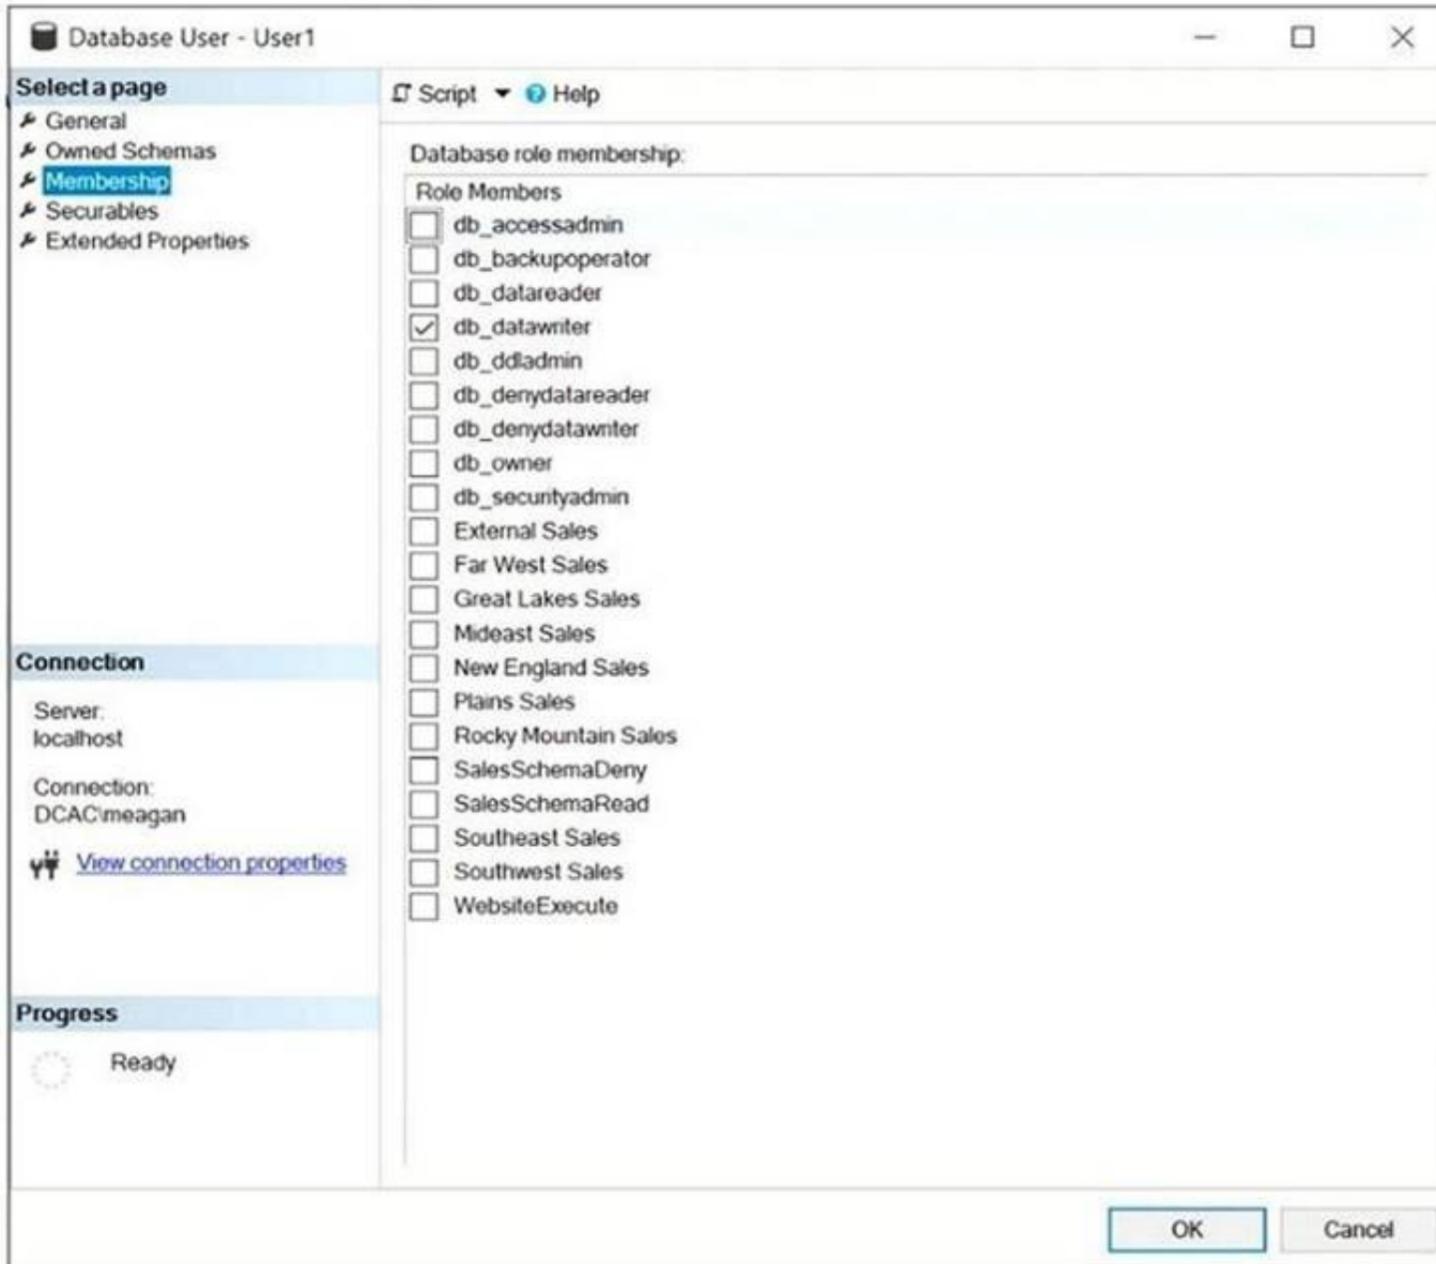
Reference:

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

NEW QUESTION 346

- (Exam Topic 5)

You have a Microsoft SQL Server database named DB1 that contains a table named Table1. The database role membership for a user named User1 is 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.

User1 can [answer choice].

	▼
add a column to Table1	
delete a row from Table1	
delete Table1	

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

	▼
db_datareader	
db_ddladmin	
db_denydatareader	
db_denydatawriter	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: delete a row from Table1

Members of the db_datawriter fixed database role can add, delete, or change data in all user tables. Box 2: db_datareader

Members of the db_datareader fixed database role can read all data from all user tables. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

NEW QUESTION 348

- (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-vcare-single-databases#business-critic>

NEW QUESTION 351

- (Exam Topic 5)

You have a burstable Azure virtual machine named VM1 that hosts an instance of Microsoft SQL Server. You need to attach an Azure ultra disk to VM1. The solution must minimize downtime on VM1.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Attach the ultra disk.
Stop and deallocate VM1.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Attach the ultra disk.
Stop and deallocate VM1.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

Answer Area

Stop and deallocate VM1.
Attach the ultra disk.
Set Enable Ultra disk compatibility to Yes .
Resize VM1.
Start VM1.

NEW QUESTION 354

- (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 358

.....

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