DP-300 Dumps

Administering Relational Databases on Microsoft Azure (beta)

https://www.certleader.com/DP-300-dumps.html

NEW QUESTION 1
- (Exam Topic 5)
You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.
You need to test impersonation of user1 in db1 by running aSELECTstatement and returning to the original execution context.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
EXECUTE AS [_____▼] = 'user1@contoso.com'
            CALLER
            LOGIN
            OWNER
            USER

GO

SELECT SUSER_SNAME()

[_____▼]
  REVERT
  REVOKE
  ROLLBACK

GO
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface Description automatically generated
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15 https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15

NEW QUESTION 2
- (Exam Topic 5)
You have an Azure subscription that uses a domain named contoso.com.
You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.
DBServer1 has a high volume of data changes and low latency requirements for data writes.
You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.
What should you do?

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

Answer: C

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

NEW QUESTION 3
- (Exam Topic 5)
You have an 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.

Run:

A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:

A Recovery Services vault
An Azure Blob storage account
An SMB file share

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 4
- (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

"dnsZonePartner"
"storageAccountType"
"SystemAssigned"
"UserAssigned"
"zoneRedundant"

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":      "SystemAssigned"
    },
    "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",

        "zoneRedundant"         : "True"
    }
}
```

NEW QUESTION 5
- (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 6
- (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 Databricksmonitoring 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 7
- (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 8
- (Exam Topic 5)
You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

CREATE USER ▼ FROM ▼

[App1]
[Contoso\app1]
[App1@contoso.com]

login
Windows
EXTERNAL PROVIDER

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 9
- (Exam Topic 5)
You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2. You need to limit the number of IOPs that App2 queries generate on SQL1.
Which two actions should you perform on SQL1? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: BD

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

NEW QUESTION 10
- (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 schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Must use an Azure Data Factory, not an Azure Databricks job. Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

NEW QUESTION 10
- (Exam Topic 5)
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:
• Capture raw event data and store the data in Azure Storage.
• Minimize the performance impact of capturing extended events.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
CREATE  EVENT SESSION session1 ON DATABASE

    ADD EVENT sqlserver.sql_statement_starting

        (

        ACTION (sqlserver.sql_text)

        WHERE statement LIKE 'UPDATE gmTabEmployee%'

        )

    ADD TARGET

        package0.    event_file            ▼

                     event_file
                     event_stream
                     ring_buffer

        )

        SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'

        )

    WITH

        (MAX_MEMORY = 10 MB,

        EVENT_RETENTION_MODE=    ALLOW_MULTIPLE_EVENT_LOSS    ▼

        MAX_DISPATCH_LATENCY = 3 SEC    ALLOW_MULTIPLE_EVENT_LOSS
                                        ALLOW_SINGLE_EVENT_LOSS
                                        NO_EVENT_LOSS
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

```
Answer Area

    CREATE  EVENT SESSION session1 ON DATABASE

      ADD EVENT sqlserver.sql_statement_starting

        (

        ACTION (sqlserver.sql_text)

        WHERE statement LIKE 'UPDATE gmTabEmployee%'

        )

    ADD TARGET

    package0.    event_file          ▼
                 event_file
                 event_stream
                 ring_buffer

        )

        SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'

        )

    WITH

        (MAX_MEMORY = 10 MB,

        EVENT_RETENTION_MODE=    ALLOW MULTIPLE EVENT LOSS          ▼
                                 ALLOW_MULTIPLE_EVENT_LOSS
        MAX_DISPATCH_LATENCY = 3 SE  ALLOW_SINGLE_EVENT_LOSS
                                 NO_EVENT_LOSS
```

NEW QUESTION 12
- (Exam Topic 5)
You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

❯ Send the output to an Azure Synapse.

❯ Identify spikes and dips in time series data.

❯ Minimize development and configuration effort.
Which should you include in the solution?

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

Answer: C

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

NEW QUESTION 14
- (Exam Topic 5)
You 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-tablesovervie

NEW QUESTION 19
- (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 geo-replication. 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 20
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies that several tables are missing indexes. You need to ensure that indexes are created for the tables.
What should you do?

A. RuntheDBCC SQLPERF command.
B. Run thedbcc dbreindexcommand.
C. Modify the automatic tuning settings for db1.
D. Modify the Query Store settings for db1.

Answer: C

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

NEW QUESTION 21
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1. You need to identify and repair any physical or logical corruption in DB1. The solution must meet the
following requirements:
• Minimize how long it takes to complete the procedure.
• Minimize data loss.
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

DBCC CHECK [DB1], [▼]) WITH [▼]

| NOINDEX |
| REPAIR_ALLOW_DATA_LOSS |
| REPAIR_FAST |
| REPAIR_REBUILD |

| EXTENDED_LOGICAL_CHECKS; |
| PHYSICAL_ONLY; |
| TABLOCK; |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

DBCC CHECK [DB1], [▼]) WITH [▼]

| NOINDEX |
| REPAIR_ALLOW_DATA_LOSS |
| REPAIR_FAST |
| REPAIR_REBUILD |

| EXTENDED_LOGICAL_CHECKS; |
| PHYSICAL_ONLY; |
| TABLOCK; |

NEW QUESTION 23
- (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 28
- (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 29
- (Exam Topic 5)
You have an Azure Databricks resource.
You need to log actions that relate to changes in compute for the Databricks resource. Which Databricks services should you log?

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

Answer: E

Explanation:
Cloud Provider Infrastructure Logs.
Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace.

Customers, especially in the regulated industries, also need records on activities like:

≫ User access control to cloud data storage

≫ Cloud Identity and Access Management roles

≫ User access to cloud network and compute

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

Reference:

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

---

NEW QUESTION 30
- (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 34
- (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 37
- (Exam Topic 5)
You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in fiveminute intervals and report only events that arrive during the interval.
The output will be sent to a Delta Lake table. Which output mode should you use?

A. complete
B. append
C. update

Answer: A

Explanation:
Complete mode: You can use Structured Streaming to replace the entire table with every batch.
Reference:
https://docs.databricks.com/delta/delta-streaming.html

NEW QUESTION 40
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You have a table name Table1 that has 20 columns of typeCHAR(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 tonvarchar(MAX).
C. Change the column type tovarchar(MAX).
D. Change the column type tovarchar(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 45
- (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 49
- (Exam Topic 5)
You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1. The only user who was created is the server administrator.
You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.
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
Connect to DB1 by using the Active Directory admin account.	
Create a user by using the FROM EXTERNAL PROVIDER clause.	
Connect to DB1 by using the server administrator account.	
Set the Active Directory Admin for AzSQL1.	
From the Azure portal, assign the SQL DB Contributor role to the user.	
Create a login in the master database.	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Set up the Active Directory Admin for AzSQL1. Step 2: Connect to DB1 by using the server administrator.
Sign into your managed instance with an Azure AD login granted with the sysadmin role. Step 3: Create a user by using the FROM EXTERNAL PROVIDER
clause.
FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level
Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:
CREATE USER [AAD_principal] FROM LOGIN [Azure AD login] Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql

NEW QUESTION 54
- (Exam Topic 5)
You have an Azure SQL managed instance.
You need to restore a database named DB1 by using Transact-SQL.
Which command should you run? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
RESTORE [ ▼ ] DB1 FROM [ ▼ ]
         DATABASE              DISK = N'\\NAS01\$SQLBackups\DB1.bak';
         FILE                  TAPE = N'\\.\Tape0'
         LOG                   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 58
- (Exam Topic 5)
You have two instances of SQL Server on Azure virtual Machines named VM1 and VM2. VM1 hosts a database named dbl.
You plan to create a database availability group (DAG) for dbl. 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		Answer Area
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.		

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions		Answer Area
From the master database, run CREATE ENDPOINT.		From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.
From db1, run CREATE CERTIFICATE.	⟩	From the master database, run CREATE CERTIFICATE.
From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.	⟨	From db1, run CREATE ENDPOINT.
From the master database, run CREATE CERTIFICATE.		
From db1, run CREATE ENDPOINT.		

NEW QUESTION 59
- (Exam Topic 5)
You have an Azure SQL database.
Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.
You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.
What should you use?

A. Performance Recommendations in the Azure portal
B. Extended Events in Microsoft SQL Server Management Studio (SSMS)
C. Query Store in Microsoft SQL Server Management Studio (SSMS)
D. Query Performance Insight in the Azure portal

Answer: C

Explanation:
Use the Query Store Page in SQL Server Management Studio.
Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.
Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query.
This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-qu

NEW QUESTION 62
- (Exam Topic 5)
You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs
Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

Actions

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

Answer Area

>
<

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

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

Answer Area

| Create a Database Mail account. |
| Create a profile named `AzureManagedInstance_dbmail_profile`. |
| Enable email notifications upon failure. |

>
<

NEW QUESTION 66
- (Exam Topic 5)
You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.
Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.
You set the Read/Write failover policy to Automatic.
What are two results of the configuration? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Answer: AE

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

NEW QUESTION 69
- (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 70
- (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 = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-

02T19:30:08Z&spr=https&sv=2021-06-

08&sr=b&sig=B%2FxEYQiOC%4BqyYCeqWHSz2QpRI%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 = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-

02T19:30:08Z&spr=https&sv=2021-06-

08&sr=b&sig=B%2FxEYQiOC%4BqyYCeqWHSz2QpRI%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 73
- (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 endpoint
D. an ExpressRoute gateway

Answer: C

Explanation:
A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.
The service could be an Azure service such as:

≫ Azure Storage

≫ Azure Cosmos DB

≫ Azure SQL Database

≫ Your own service using a Private Link Service. Reference:
https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview

NEW QUESTION 77
- (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 79
- (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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

| 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. |

Answer Area

| From the Azure portal, export the Azure Resource Manager templates |
| Change the server name and related variables in the templates |
| From the Azure portal, deploy the templates. |
| From the database project, deploy the database schema and permissions |

NEW QUESTION 80
- (Exam Topic 5)
You have 10 Azure virtual machines that have SQL Server installed.
You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups.
What should you include in the backup strategy?

A. Automated Backup in the SQL virtual machine settings
B. Azure Backup
C. Azure Site Recovery
D. SQL Server Agent jobs

Answer: B

Explanation:
Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault.
There are several advantages that this solution provides, especially for Enterprises.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/backup-restore#azbackup

NEW QUESTION 82
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named
container1.
You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

NEW QUESTION 85
- (Exam Topic 5)
You have a Microsoft SQL Server 2017 server that hosts five databases. You Plan to migrate the databases to Azure.
You need to recommend a solution that meets the following requirements:
‚ Automatically scales compute based on the workload demand
‚ Provides per-second billing
What should you include in the recommendation? To answer, select the appropriate options in the answer area.

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard
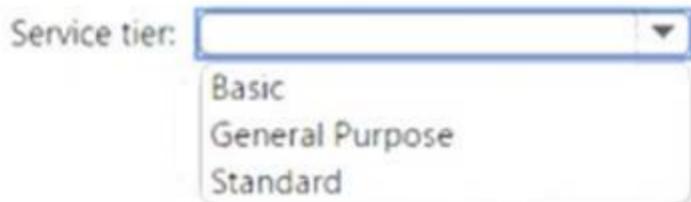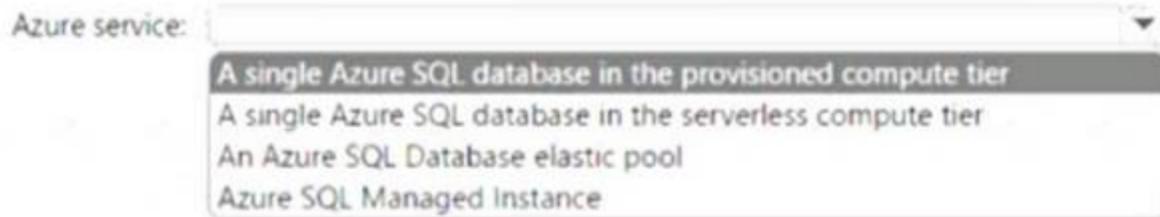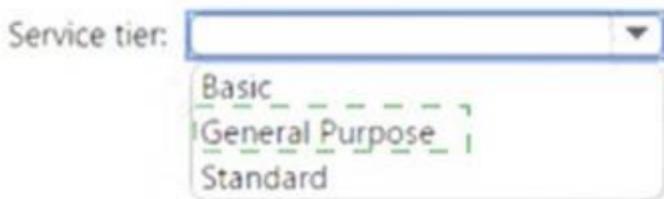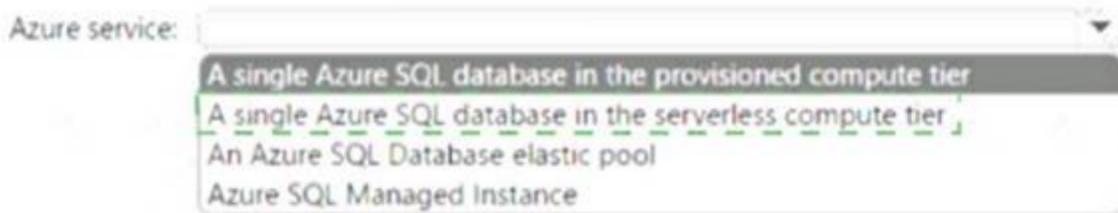
NEW QUESTION 86
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1.
What should you do?

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

Answer: D

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

NEW QUESTION 89
- (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 90
- (Exam Topic 5)
You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database. A recent failure causes the master database to be lost.
You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual
machine.
Which four actions should you perform in sequence to resolve the issue? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct.

Actions

| Add a certificate to an Azure key vault |
| Enable Transparent Data Encryption (TDE) |
| Encrypt a copy of the master key by using the service master key |
| Turn on the TRUSTWORTHY property and the CLR property |
| Attach the SSISDB database |
| Open the master key for the SSISDB database |

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Attach the SSISDB database
Step 2: Turn on the TRUSTWORTHY property and the CLR property
If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr)
Step 3: Open the master key for the SSISDB database
Restore the master key by this method if you have the original password that was used to create SSISDB.
open master key decryption by password = 'LS1Setup!' --'Password used when creating SSISDB' Alter Master Key Add encryption by Service Master Key
Step 4: Encrypt a copy of the mater key by using the service master key Reference:
https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog

NEW QUESTION 95
- (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 vCPUSs 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 98
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.
Which deployment option and pricing tier should you configure?

A. Azure SQL Database Hyperscale

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

Answer: C

Explanation:
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

NEW QUESTION 102
- (Exam Topic 5)
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver

NEW QUESTION 107
- (Exam Topic 5)
You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.
You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:
The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
Administrative effort must be minimized. What should you use to host the databases?

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

Answer: D

NEW QUESTION 109
- (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.

```
[ ▼ ]  DATABASE MyDB1
┌─────────────────┐
│ BACKUP          │
│ CREATE          │
│ RESTORE         │
└─────────────────┘
     FROM DISK = 'D:\DB1.bak'

     WITH  [        ▼ ]
GO        ┌─────────────────┐
          │ NORECOVERY      │
          │ RECOVERY        │
          │ STANDBY         │
          └─────────────────┘
```

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
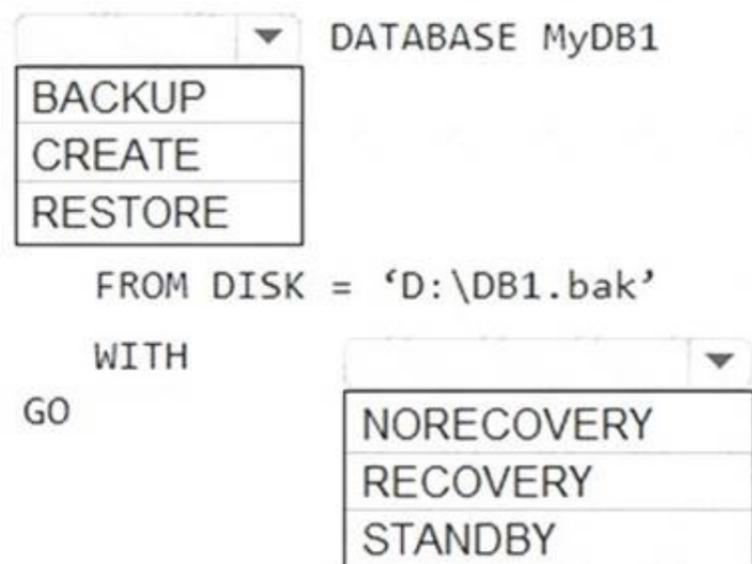
NEW QUESTION 110
- (Exam Topic 5)
You have a new Azure subscription.
You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1. You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that 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 112
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database named SQL1. SQL1 is in an Azure region that does not support availability zones.
You need to ensure that you have a secondary replica of SQLI in the same region. What should you use?

A. log shipping
B. auto-failover groups
C. active geo-replication
D. Microsoft SQL Server failover clusters

Answer: C

NEW QUESTION 113
- (Exam Topic 5)
You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD sen/ice principal named appl.
You need to ensure that appl can access db1. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To

answer, select the appropriate options in the
answer area.
NOTE: Each correct selection is worth one point.

Answer Area

CREATE [▼] [app1] [▼]

CREDENTIAL		FOR LOGIN app1
LOGIN		FROM EXTERNAL PROVIDER
USER		FROM LOGIN app1
		WITHOUT LOGIN

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

CREATE [▼] [app1] [▼]

CREDENTIAL		FOR LOGIN app1
LOGIN		**FROM EXTERNAL PROVIDER**
USER		FROM LOGIN app1
		WITHOUT LOGIN

NEW QUESTION 114
- (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 117
- (Exam Topic 5)
You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases.
Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: ACE

Explanation:
CE: Estimate the vCores needed for the pool as follows:
For vCore-based purchasing model: MAX(<Total number of DBs X average vCore utilization per DB>,
<Number of concurrently peaking DBs X Peak vCore utilization per DB)
A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

NEW QUESTION 122
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1. You need to move a SQL Server Agent job from DB1 to SQLMI1. Which job attribute is unsupported in SQLMI1?

A. log to table
B. email notifications
C. schedules
D. output files

Answer: D

NEW QUESTION 126
- (Exam Topic 5)
You have an Azure subscription.
You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.
You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs. The solution must minimize administrative effort.
Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

MSX:

SQL database
SQL managed instances
SQL virtual machines

TSX:

SQL database
SQL managed instances
SQL virtual machines

NEW QUESTION 130
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.
You need to resolve the performance issues. What should you do?

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

Answer: A

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

NEW QUESTION 133
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains a table named Table 1. You run a query to load data into Table1.
The performance metrics of Table1 during the load operation are shown in the following exhibit.
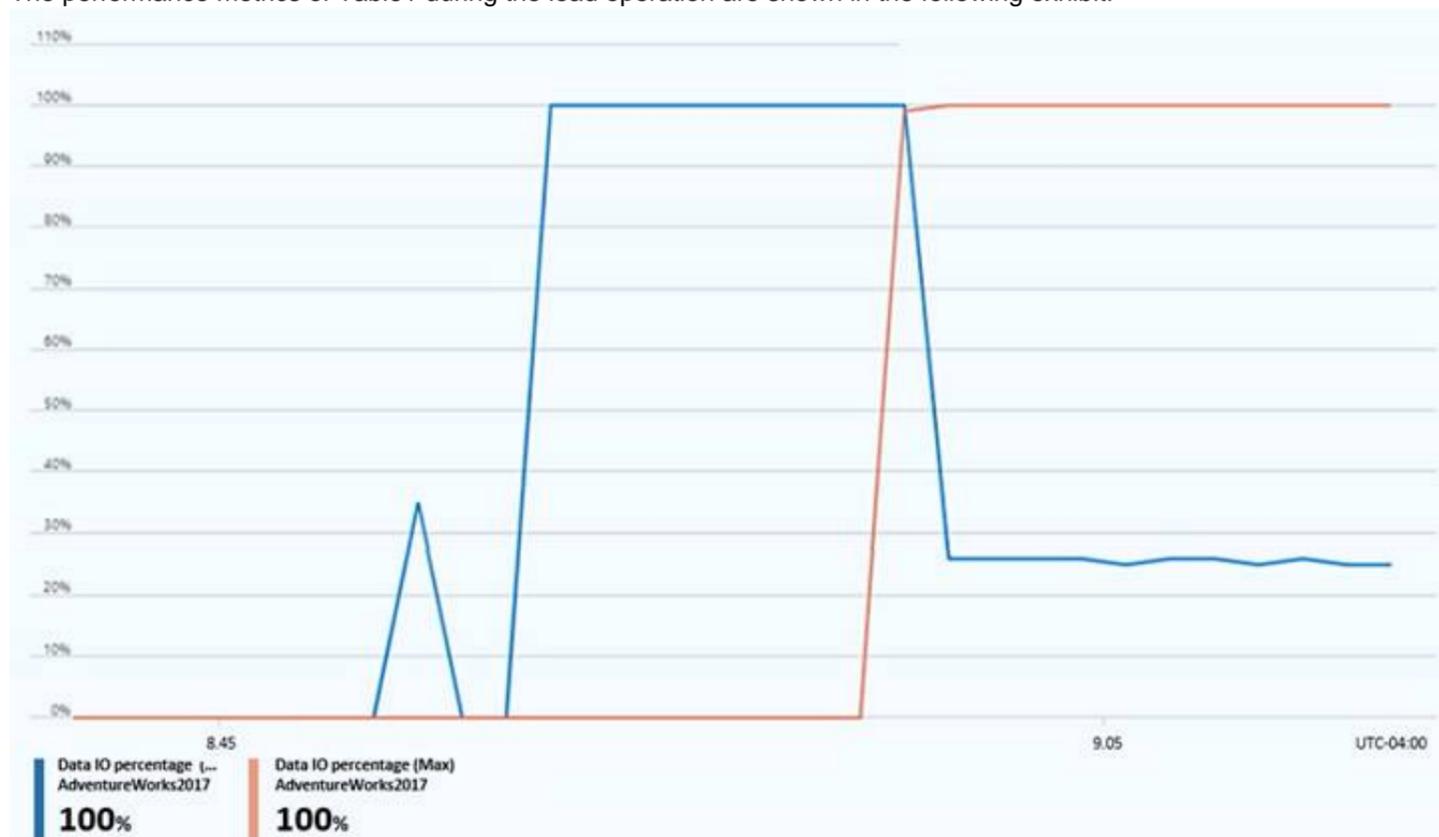
Data IO percentage (... AdventureWorks2017	Data IO percentage (Max) AdventureWorks2017
100%	**100%**

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

▼
scale the resource
use an elastic resource
perform query tuning

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

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

▼
~~scale the resource~~
use an elastic resource
perform query tuning

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

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

NEW QUESTION 135
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains a private certificate named Sales. The private key for Sales is encrypted with a password. You need to change the password for the private key. Which
Transact-SQL statement should you run?
A)

```
ALTER CERTIFICATE Sales
    WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%',
    ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```

B)

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

C)

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

D)

```
ALTER CERTIFICATE Sales    WITH PRIVATE KEY (DECRYPTION BY PASSWORD = ' EWYx9Xk+$#');
```

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

Answer: C

NEW QUESTION 136
- (Exam Topic 5)
You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.
You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.
Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

MFA:

Azure AD authentication
Microsoft SQL Server authentication
Passwordless authentication
Windows authentication

Database-level authentication:

Application roles
Contained database users
Database roles
Microsoft SQL Server logins

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, chat or text message Description automatically generated
Box 1: Azure AD authentication
Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.
Box 2: Contained database users
Azure Active Directory Uses contained database users to authenticate identities at the database level. Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-authentication

NEW QUESTION 137
- (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 138
- (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 139
- (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 140
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.
You need to protect the contents of Column1 at rest, in transit, and in use.
How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Encryption key:
- Column encryption key
- Database encryption key
- Service master key

Encryption type:
- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Column encryption Key
Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine

NEW QUESTION 145
- (Exam Topic 5)
You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.
You need to implement alerts by using Azure Monitor. The solution must meet the following requirements: ≫ Minimize costs.
≫ Aggregate Intelligent Insights telemetry from each database. What should you do?

A. From the Diagnostic settings of each database, select Send to Log Analytics.
B. From the Diagnostic settings of each database, select Stream to an event hub.
C. From the Diagnostic settings of SQLMI1. select Send to Log Analytics.
D. From the Diagnostic settings of SQLMI1. select Stream to an event hub.

Answer: A

Explanation:
Reference:

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

NEW QUESTION 147
- (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 150
- (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 152
- (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 an Azure Databricks notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 156
- (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.

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 159
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named
container1.
You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead use a serverless SQL pool to create an external table with the extra column. Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

NEW QUESTION 161
- (Exam Topic 5)
You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

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

You plan to enable Always Encrypted for Table1.
Which two columns support encryption? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point

A. Column1
B. Column2
C. Column3
D. Column4
E. Column5

Answer: AD

NEW QUESTION 164
- (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 167
- (Exam Topic 5)
You have the following Azure Resource Manager template.

```
...
    "variable": {
        "serverName": "azsqldbserver0001"
    },
    "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"
            },
            "resources": [
                {
                    "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.	○	○
The template deploys a database to an Azure SQL Database managed instance.	○	○
The pricing tier of the database deployment is based on DTUs.	○	○

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 171
- (Exam Topic 5)
You receive numerous alerts from Azure Monitor for an Azure SQL database.
You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Set Threshold Sensitivity to High
B. Set the Alert logic threshold to Dynamic
C. Set the Alert logic threshold to Static
D. Set Threshold Sensitivity to Low
E. Set Force Plan to On

Answer: BD

Explanation:
B: Dynamic Thresholds continuously learns the data of the metric series and tries to model it using a set of algorithms and methods. It detects patterns in the data such as seasonality (Hourly / Daily / Weekly), and is able to handle noisy metrics (such as machine CPU or memory) as well as metrics with low dispersion (such as availability and error rate).
D: Alert threshold sensitivity is a high-level concept that controls the amount of deviation from metric behavior required to trigger an alert.
Low – The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds

NEW QUESTION 174
- (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 176
- (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 Appl as soon as possible, without considering immediate costs
What you do?

A. Increase the number of vCores allocated to db1.
B. Decrease the auto-pause delay for db1.
C. Disable auto-pause delay for db1.
D. Enable automatic tuning for db1.

Answer: D

NEW QUESTION 179
- (Exam Topic 5)
You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1. You need to configure SQL! to use mixed mode authentication. Which procedure should you run?

A. sp_addremotelogin
B. xp_instance_regwrite
C. sp_cnarge_users_login
D. xp_grant_login

Answer: B

NEW QUESTION 180
- (Exam Topic 5)
You are building an Azure Stream Analytics job to retrieve game data.
You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.
How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
SELECT [                                      ▼] as HighestScore
       | Collect(Score)                        |
       | CollectTop(1)OVER(ORDER BY Score Desc) |
       | Game, MAX(Score)                       |
       | TopOne() OVER(PARTITION BY Game ORDER BY Score Desc) |

FROM input TIMESTAMP BY CreatedAt

GROUP BY [                                  ▼]
         | Game                                              |
         | Hopping(minute, 5)                                |
         | Tumbling(minute, 5)                               |
         | Windows(TumblingWindow(minute, 5), Hopping(minute, 5)) |
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)
TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.
Analytic Function Syntax:
TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT
DURATION clause> [<WHEN clause>])
Box 2: Tumbling(minute 5)
Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

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

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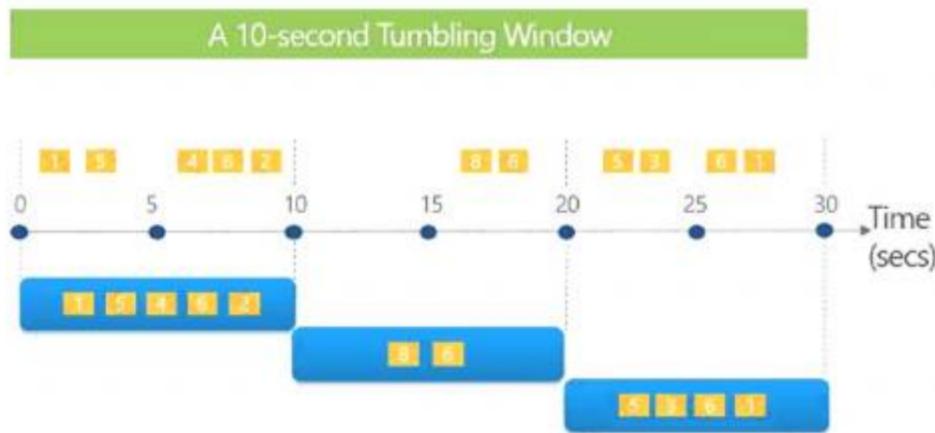
```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

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

NEW QUESTION 184
- (Exam Topic 5)
You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.
Which switch should you use to switch between languages?

A. \\[<language>]
B. %<language>
C. \\[<language>]
D. @<language>

Answer: B

Explanation:

You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are:
%python, %r, %scala, and %sql.

Reference:
https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use

NEW QUESTION 188
- (Exam Topic 5)
You have an Azure SQL database.
You have a query and the associated execution plan as shown in the following exhibit.
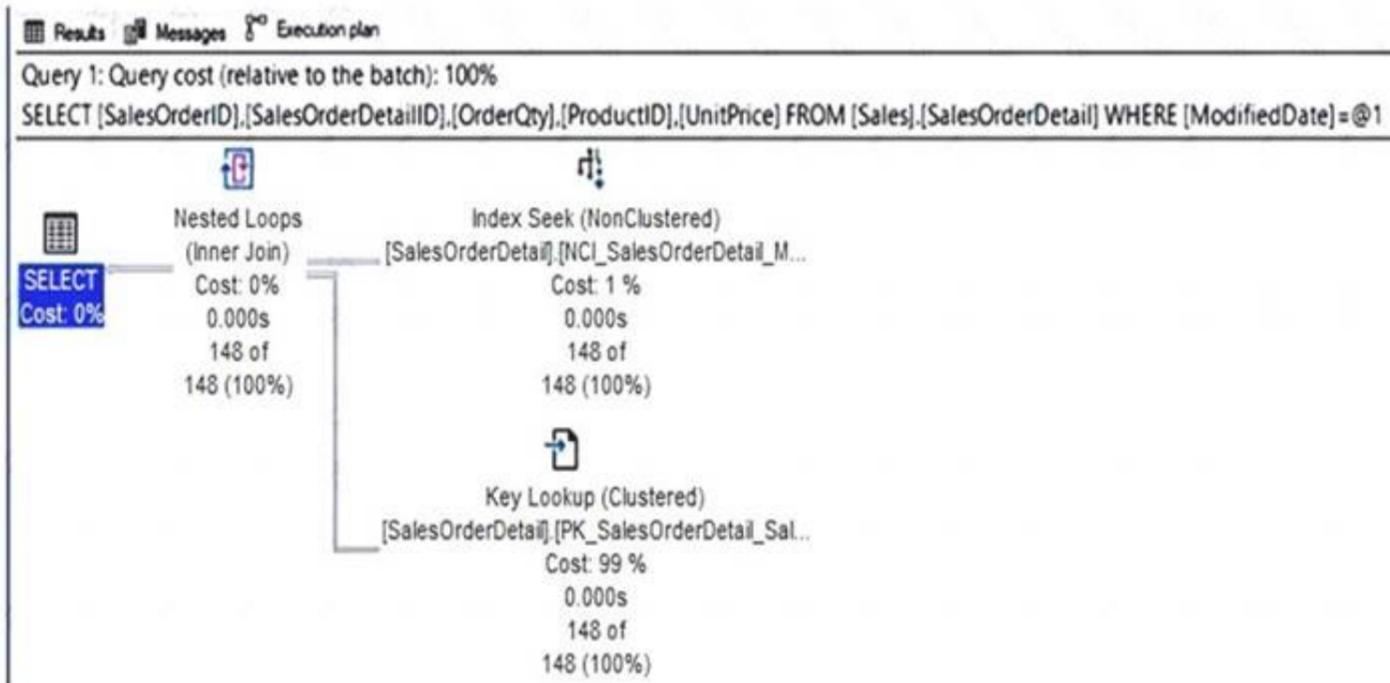
Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE:Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Box 1: Key Lookup
The Key Lookup cost is 99% so that is the performance bottleneck. Box 2: nonclustered index
The key lookup on the clustered index is used because the nonclustered index does not include the required columns to resolve the query. If you add the required columns to the nonclustered index, the key lookup will not be required.

NEW QUESTION 193
- (Exam Topic 5)
You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.
You plan to create a database named DB1 in Pool1.
You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.
Which format should you use for the tables in DB1?

A. JSON
B. CSV
C. Parquet
D. ORC

Answer: C

Explanation:
Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.
For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can

shut down your Spark pools and still query Spark external tables from serverless SQL pool.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables

NEW QUESTION 198
- (Exam Topic 5) You have an Azure SQL database. You identify a long running query.
You need to identify which operation in the query is causing the performance issue.
What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

A. Live Query Statistics
B. an estimated execution plan
C. an actual execution plan
D. Client Statistics

Answer: A

Explanation:
https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/

NEW QUESTION 200
- (Exam Topic 5)
You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.
You execute the Transact-SQL query shown in the following exhibit.

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

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

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

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

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

a header
a data
an empty

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 203
- (Exam Topic 5)
You have a SQL pool in Azure Synapse that contains a table named dbo.Customers. The table contains a column name Email.
You need to prevent nonadministrative users from seeing the full email addresses in the Email column. The users must see values in a format of
aXXX@XXXX.com instead.
What should you do?

A. From the Azure portal, set a mask on the Email column.
B. From the Azure portal, set a sensitivity classification of Confidential for the Email column.
C. From Microsoft SQL Server Management Studio, set an email mask on the Email column.
D. From Microsoft SQL Server Management Studio, grant the SELECT permission to the users for all the columns in the dbo.Customers table except Email.

Answer: B

Explanation:
The Email masking method, which exposes the first letter and replaces the domain with XXX.com using a constant string prefix in the form of an email address.
Example: aXX@XXXX.com

NEW QUESTION 206
- (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 209
- (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 210
- (Exam Topic 5)
You have an Azure subscription.
You plan to deploy an Azure SQL database by using an Azure Resource Manager template.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Answer Area

```
{
  "resources": [
    {
      "type": [                                      ▼ ]
             "Microsoft.Sql/servers"
             "Microsoft.SqlVirtualMachines/sqlVirtualMachines"
             "Microsoft.Synapse/workspaces/sqldatabases"

      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",

          ...
          [         ▼ ]
          "dependsOn":[
          "properties":[
          "tags":[

          "[resourceId('Microsoft.Sql/servers',concat(parameters('name1')))]"

          ...
      }
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart

NEW QUESTION 212
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to display the estimated execution plan of a query by using the query editor in the Azure portal.
What should you do first?

A. Run theset showplan_allTransact-SQL statement.
B. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
C. Run theset forceplanTransact-SQL statement.
D. Enable Query Store for DB1.

Answer: A

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

NEW QUESTION 215
- (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 218
- (Exam Topic 5)
You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.
You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text, table Description automatically generated
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-def

NEW QUESTION 219
- (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 222
- (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 223
- (Exam Topic 5)
You are developing an application that uses Azure Data Lake Storage Gen 2.
You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Answer: C

Explanation:
A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:
What resources the client may access.
What permissions they have to those resources. How long the SAS is valid.
Note: Data Lake Storage Gen2 supports the following authorization mechanisms:
➢ Shared Key authorization
➢ Shared access signature (SAS) authorization
➢ Role-based access control (Azure RBAC)
➢ Shared Key authorization
➢ Shared access signature (SAS) authorization
➢ Role-based access control (Azure RBAC)
➢ Access control lists (ACL)
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 228
- (Exam Topic 5)
You have an Azure SQL database.
You run the following PowerShell script.

```
$serverName = "SERVER1"
$resourceGroup = "RG1"
$dbName = "DB1"

Connect-AzAccount

$server = Get-AzSqlServer -ServerName $serverName -ResourceGroupName
$resourceGroup

Set-AzSqlDatabaseBackupShortTermRetentionPolicy -ResourceGroupName $resourceGroup
-ServerName $server '
    -DatabaseName $dbName -RetentionDays 21

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $serverName -
DatabaseName $dbName '
    -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention PSY
-WeekOfYear 52
```

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
DB1 can be restored to a specific point in time 30 days ago.	○	○
DB1 can be restored from a weekly backup performed six months ago.	○	○
DB1 can be restored from a yearly backup performed six years ago.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackupshorttermretentionpolicy?vi https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackuplongtermretentionpolicy?vie

NEW QUESTION 232
- (Exam Topic 5)
You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Duplicate Azure Policy definitions

Run Azure Policy remediation tasks

Create an Azure Blueprints assignment

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create an Azure Policy Initiative
The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.
With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.
Step 2: Create an Azure Policy Initiative assignment
Assign the initiative definition you created in the previous step. Step 3: Run Azure Policy remediation tasks
To apply the Policy Initiative to the existing SQL databases. Reference:
https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

NEW QUESTION 236
- (Exam Topic 5)
You have an Azure SQL database named sqldb1.
You need to minimize the amount of space by the data and log files of sqldb1. What should you run?

A. DBCC SHRINKDATABASE
B. sp_clean_db_free_space
C. sp_clean_db_file_free_space
D. DBCC SHRINKFILE

Answer: D

Explanation:
DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database. Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql

NEW QUESTION 239
- (Exam Topic 5)
You have an Azure SQL database named DB3.
You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click theExhibittab.)
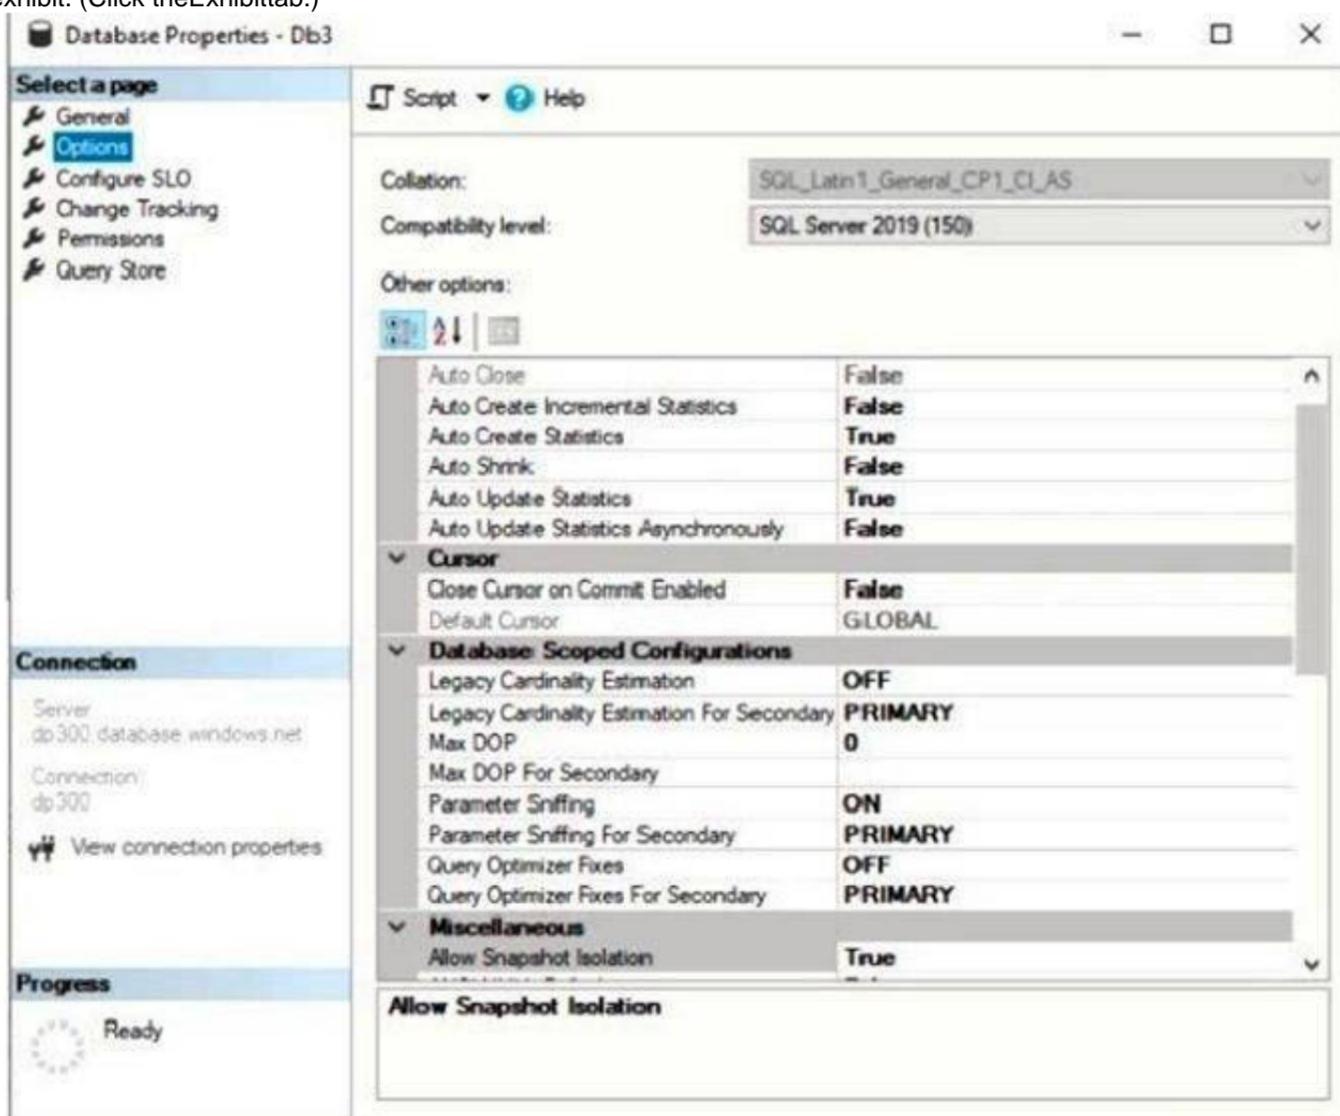
Which Transact-SQL command should you run?

A. GRANT SHOWPLAN TO DevUser
B. GRANT VIEW DEFINITION TO DevUser
C. GRANT VIEW DATABASE STATE TO DevUser
D. GRANT SELECT TO DevUser

Answer: C

Explanation:
The exhibits displays Database [State] properties.
To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-properties-options-page

NEW QUESTION 243
- (Exam Topic 5)
You have an Azure SQL logical server. You run the following script.

```
CREATE DATABASE Sales
GO
CREATE TABLE [dbo].[Orders]
  (
     [OrderID] INT NOT NULL,
     [OrderDescription] NVARCHAR (MAX) NOT NULL,
     [Timestamp] Datetime2 NOT NULL
  )
  WITH (
     SYSTEM_VERSIONING = ON,
     LEDGER = ON
     );
GO
```

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 Orders table will allow only rows to be inserted.	○	○
To create additional tables in the Sales database, the LEDGER = ON parameter must be used.	○	○
To ensure that a timestamp is added to each row in the Orders table, the GENERATED ALWAYS	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Yes
No No

NEW QUESTION 246
- (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 248
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:

> During normal operations, provide at least two readable copies of Sales.

> Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

NEW QUESTION 249
- (Exam Topic 5)
You have an Azure SQL database named DB1. A user named User 1 has an Azure AD account.
You need to provide User1 with the ability to add and remove columns from the tables inDBV The solution must use the principle of least privilege.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point

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

Answer: AD

NEW QUESTION 252
- (Exam Topic 5)
You have an Azure subscription.
You plan to deploy an instance of SQL Sevier on Azure Virtual Machines that supports Write Accelerator.
Which virtual machine series should you use?

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

Answer: C

NEW QUESTION 256
- (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.

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 257
- (Exam Topic 5)
Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.
You plan to migrate db1 to an Azure SQL managed instance.
You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

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

Answer: D

NEW QUESTION 258
- (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 260
- (Exam Topic 5)
You have a on-premises Microsoft SQL Server named SQL1 that hosts five databases.
You need to migrate the databases to an Azure SQL managed instance. The solution must minimize downtime and prevent data loss.
What should you use?

A. log shipping
B. Always On availability groups
C. Database Migration Assistant
D. Backup and Restore

Answer: C

NEW QUESTION 261
- (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 264
- (Exam Topic 5)
You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:
• Support real-time data replication to a different geographic region.
• Use Azure as a disaster recovery target.
• Minimize costs and administrative effort.
What should you include in the recommendation?

A. database mirroring on an instance of SQL Server on Azure Virtual Machines
B. availability groups for SQL Server on Azure Virtual Machines
C. an Azure SQL Managed Instance link
D. transactional replication to an Azure SQL Managed Instance

Answer: D

NEW QUESTION 269
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1. You need to configure the autogrow and autoshrink settings for DB1.
Which statements should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Autogrow:

| ALTER DATABASE MODIFY File and Filegroup options |
| ALTER DATABASE SCOPED CONFIGURATION |
| ALTER DATABASE SET options |

Autoshrink:

| ALTER DATABASE MODIFY File and Filegroup options |
| ALTER DATABASE SCOPED CONFIGURATION |
| ALTER DATABASE SET options |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
https://learn.microsoft.com/en-us/troubleshoot/sql/admin/considerations-autogrow-autoshrink

NEW QUESTION 273
- (Exam Topic 5)
You have an Azure AD tenant and a logical Microsoft SQL server named SQL1 that hosts several Azure SQL databases.

You plan to assign Azure AD users permissions to the databases automatically by using Azure Automation. You need to create the required Automation accounts. Which two accounts should you create? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. From the Azure Active Directory admin center, create a service principal.
B. From the Azure Active Directory admin center, create a user-assigned managed identity for SQL1.
C. On SQL1. create a SQL user in the databases.
D. On SQL1, create a SQL login.
E. From the Azure Active Directory admin center, create an external identity.

Answer: AC

NEW QUESTION 275
- (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   [ ▼ ]   WITH ROLLBACK IMMEDIATE;
GO
          OFFLINE
          ONLINE
          SINGLE_USER
          TRUSTWORTHY

DBCC CHECKDB ('DB1',   [ ▼ ]   WITH NO_INFOMSGS;
GO
          MOINDEX
          PHYSICAL_ONLY
          REPAIR_ALLOW_DATA_LOSS
          REPAIR_FAST

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 278
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.
You plan to migrate DB1 to an Azure SQL Database managed instance.
What should you use to minimize downtime and data loss during the migration?

A. database mirroring
B. distributed availability groups
C. Always On Availability Group
D. Azure Database Migration Service

Answer: D

NEW QUESTION 281
- (Exam Topic 5)
You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.
You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

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

D. Create a cluster policy in workspace1.

Answer: C

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

NEW QUESTION 283
- (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

| 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. |

Answer Area

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 288
- (Exam Topic 5)
You have 50 Azure SQL databases.
You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.
What should you do?

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

Answer: D

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

NEW QUESTION 292

- (Exam Topic 5)
You have an Azure SQL database named DBI that contains a nonclustered index named index1. End users report slow queries when they use index1.
You need to identify the operations that are being performed on the index. Which dynamic management view should you use?

A. `sys.dm_exec_query_plan_stats`

B. `Sys.dm_db_index_physical_stats`

C. `Sys.dm_db_index_operational_stats`

D. `Sys.dm_db_index_useage_stats`

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

Answer: D

NEW QUESTION 293
- (Exam Topic 5)
You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1. You need to encrypt the data in Col1.
Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Create a database master key.

Create a column master key.

Open the symmetric key.

Create a certificate.

Update Col1.

Create a symmetric key.

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Table Description automatically generated
Reference:
https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/

NEW QUESTION 295
- (Exam Topic 5)
You have an Azure SQL database named sqldb1.
You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

A. Halve the value of Data Flush Interval.
B. Double the value of Statistics Collection Interval.
C. Halve the value of Statistics Collection Interval.
D. Double the value of Data Flush Interval.

Answer: A

Explanation:

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the
Maximum size limit between storage size checks, it transitions to read-only mode. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store

NEW QUESTION 299
- (Exam Topic 5)
You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.
You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.
You need to recommend a database service for the deployment. The solution must minimize administrative effort.
What should you include in the recommendation?

A. Azure SQL Managed instance
B. Azure SQL Database single database
C. Azure Database for PostgreSQL
D. SQL Server on Azure virtual machines

Answer: B

Explanation:
Azure SQL Database single database supports Data Sync. Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison

NEW QUESTION 301
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL managed instance that meets the following requirements:
•Optimize latency.
•Maximize the memory-to-vCore ratio.
Which service tier and hardware generation should you use? To answer, select the apocopate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Service tier:	Business Critical ▼
	Business Critical
	General Purpose
	Hyperscale
Hardware generation:	Premium-series - memory optimized ▼
	Premium-series
	Standard-series (Gen 5)
	Premium-series - memory optimized

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Service tier:	Business Critical ▼
	Business Critical
	General Purpose
	Hyperscale
Hardware generation:	Premium-series - memory optimized ▼
	Premium-series
	Standard-series (Gen 5)
	Premium-series - memory optimized
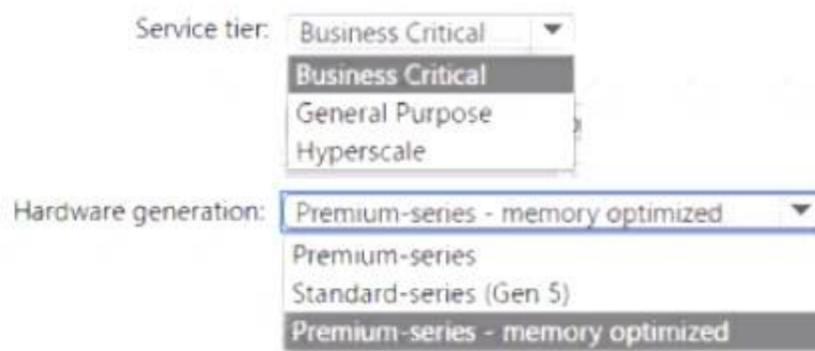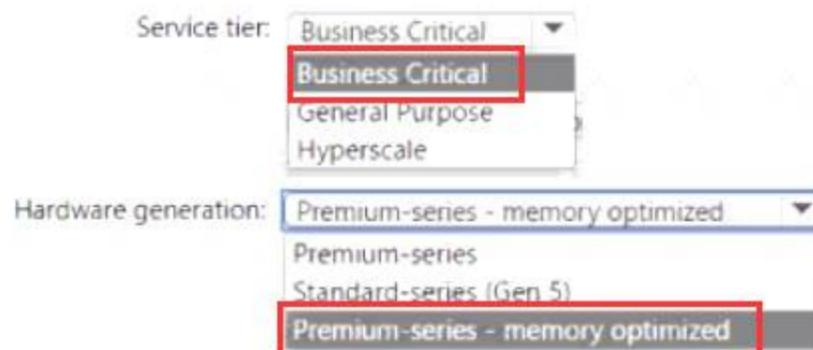
NEW QUESTION 304
- (Exam Topic 5)
You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.
Data to be loaded is identified by a column named LastUpdatedDate in the source table. You plan to execute the pipeline every four hours.
You need to ensure that the pipeline execution meets the following requirements:
Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.
Which type of trigger should you use?

A. tumbling window
B. on-demand
C. event
D. schedule

Answer: A

Explanation:
The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers

NEW QUESTION 306
- (Exam Topic 4)
You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Table type to store retail store data:

| Hash |
| Replicated |
| Round-robin |

Table type to store promotional data:

| Hash |
| Replicated |
| Round-robin |

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: Hash Scenario:
Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.
A hash distributed table can deliver the highest query performance for joins and aggregations on large tables. Box 2: Round-robin
Scenario:
You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.
A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:
» When you cannot identify a single key to distribute your data.
» If your data doesn't frequently join with data from other tables.
» When there are no obvious keys to join.
Reference:
https://rajanieshkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/

NEW QUESTION 311
- (Exam Topic 4)
You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements.
What should you create?

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

Answer: B

Explanation:
Scenario: Contoso requirements for the sales transaction dataset include: Implement a surrogate key to account for changes to the retail store addresses.
A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesidentity

NEW QUESTION 316
- (Exam Topic 2)
You are evaluating the role assignments.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE:Each correct selection is worth one point.

Statements	Yes	No
DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.	○	○
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	○	○
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	○	○

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 318
- (Exam Topic 2)
What should you implement to meet the disaster recovery requirements for the PaaS solution?

A. Availability Zones
B. failover groups
C. Always On availability groups
D. geo-replication

Answer: B

Explanation:
Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.
The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

NEW QUESTION 323
- (Exam Topic 1)
You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Purchasing model:
| Azure virtual machine reserved instances |
| DTU |
| vCore |

Deployment option:
| An Azure SQL Database elastic pool |
| An Azure SQL Database managed instance |
| A SQL Server Always On availability group |

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Box 1: DTU
Scenario:

≫ The 30 new databases must scale automatically.

≫ Once all requirements are met, minimize costs whenever possible.
You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.
In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.
Box 2: An Azure SQL database elastic pool
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

NEW QUESTION 325
- (Exam Topic 1)
You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.
What should you include in the migration plan? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Azure Database Migration Service pricing tier: ▼
Standard 2-vCore
Standard 4-vCore
Premium 4-vCore

Required Azure resource: ▼
A virtual network that has service endpoints
A VPN gateway
An Azure Logic app

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Azure Database Migration service Box 1: Premium 4-VCore
Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.

≫ Minimize downtime during the migration of the SERVER1 databases.
Premimum 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.
Reference: https://azure.microsoft.com/pricing/details/database-migration/
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online

NEW QUESTION 330
- (Exam Topic 1)
You create all of the tables and views for ResearchDB1.
You need to implement security for ResearchDB1. The solution must meet the security and compliance requirements.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=az

NEW QUESTION 332
- (Exam Topic 1)
You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals.
What should you include in the recommendation?

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

Answer: B

NEW QUESTION 334
- (Exam Topic 1)
You need to identify the cause of the performance issues on SalesSQLDb1.
Which two dynamic management views should you use? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. sys.dm_pdw_nodes_tran_locks
B. sys.dm_exec_compute_node_errors
C. sys.dm_exec_requests
D. sys.dm_cdc_errors
E. sys.dm_pdw_nodes_os_wait_stats
F. sys.dm_tran_locks

Answer: AE

Explanation:
SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.
A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.
E: Example:
The following query will show blocking information. SELECT
t1.resource_type, t1.resource_database_id, t1.resource_associated_entity_id, t1.request_mode, t1.request_session_id, t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
Note: Depending on the system you're working with you can access these wait statistics from one of three locations:
sys.dm_os_wait_stats: for SQL Server sys.dm_db_wait_stats: for Azure SQL Database
sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-lock

NEW QUESTION 339
......

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your DP-300 Exam with Our Prep Materials Via below:

https://www.certleader.com/DP-300-dumps.html