Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

NEW QUESTION 1
- (Exam Topic 5)
You have a new Azure SQL database. The database contains a column that stores confidential information. You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.
Which three actions should you perform? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

A. Turn on auditing and write audit logs to an Azure Storage account.
B. Add extended properties to the column.
C. Turn on Advanced Data Security for the Azure SQL server.
D. Apply sensitivity labels named Highly Confidential to the column.
E. Turn on Azure Advanced Threat Protection (ATP).

Answer: ACD

Explanation:
C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities. ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data
D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced,
sensitivity-based auditing and protection scenarios.
A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data_sensitivity_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview

NEW QUESTION 2
- (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()
[▼]
GO
```
REVERT
REVOKE
ROLLBACK

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 3
- (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 4
- (Exam Topic 5)
You have an Azure SQL database.
You discover that the plan cache is full of compiled plans that were used only once.
You run theselect * from sys.database_scoped_configurationsTransact-SQL command and receive the results shown in the following table.

configuration_id	name	value	is_value_default
1	LEGACY_CARDINALITY_ESTIMATION	0	1
2	QUERY_OPTIMIZER_HOTFIXES	0	1
3	OPTIMIZE_FOR_AD_HOC_WORKLOADS	0	1
4	ACCELERATED_PLAN_FORCING	1	1

You need relieve the memory pressure. What should you configure?

A. LEGACY_CARDINALITY_ESTIMATION
B. QUERY_OPTIMIZER_HOTFIXES
C. OPTIMIZE_FOR_AD_HOC_WORKLOADS
D. ACCELERATED_PLAN_FORCING

Answer: C

Explanation:
OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }
Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql

NEW QUESTION 5
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.
You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.
You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.
What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

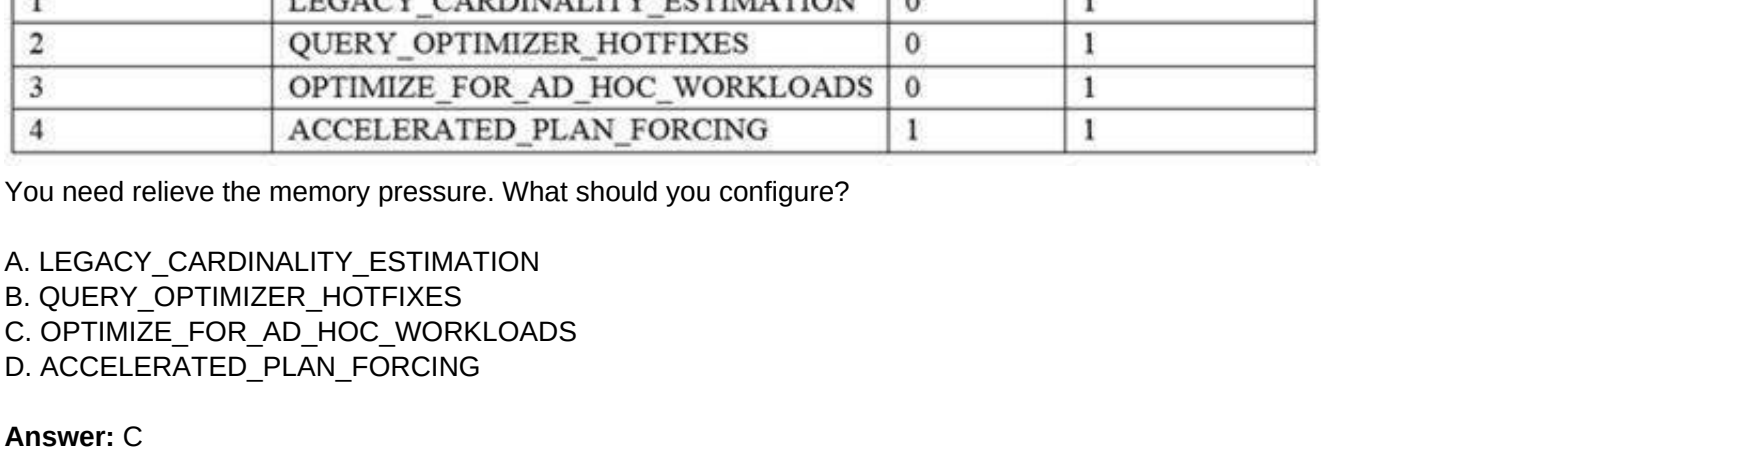

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, application, Word Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-managed-instance-online

NEW QUESTION 6
- (Exam Topic 5)
You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.
You need to ensure that SOL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

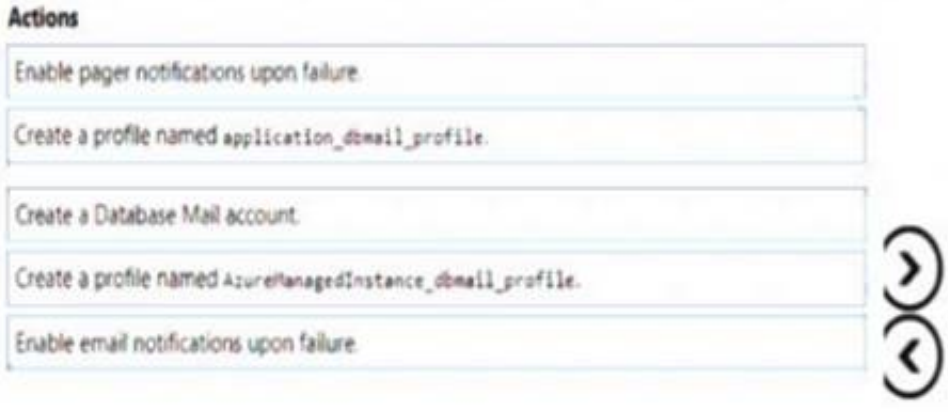

A. Mastered
B. Not Mastered

Answer: A

Explanation:

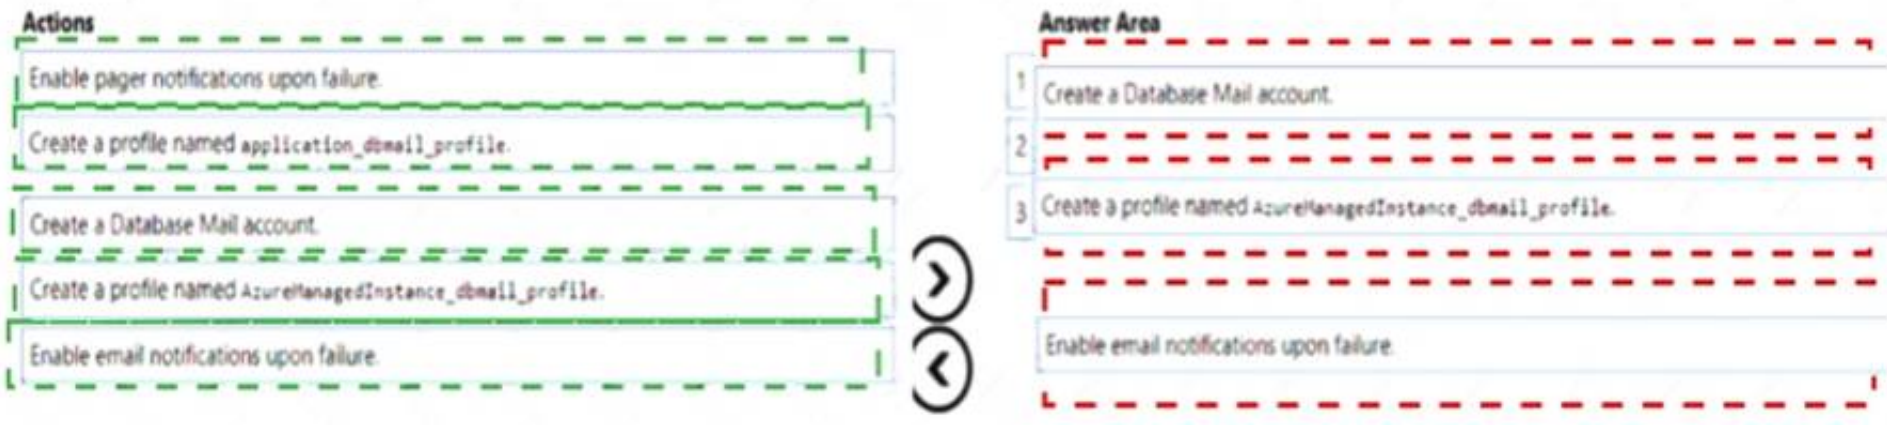

NEW QUESTION 7
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL database. The solution must meet the following requirements:
• Dynamically scale CPU resources.
• Ensure that the database can be paused to reduce costs. What should you use?

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

A. an annotation
B. a resource tag
C. a run group ID
D. a user property
E. a correlation ID

Answer: A

Explanation:
Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.
Reference:
https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/

NEW QUESTION 9
- (Exam Topic 5)
You have a Microsoft SQL Server 2017 server.
You need to migrate the server to Azure. The solution must meet the following requirements:
• Ensure that the latest version of SQL Server is used.
• Support the SQL Server Agent service. Minimize administrative effort.
What should you use?

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

Answer: D

NEW QUESTION 10
- (Exam Topic 5)
You have an Azure SQL database named DB 1 in the General Purpose service tier. You need to monitor DB 1 by using SQL Insights.
What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

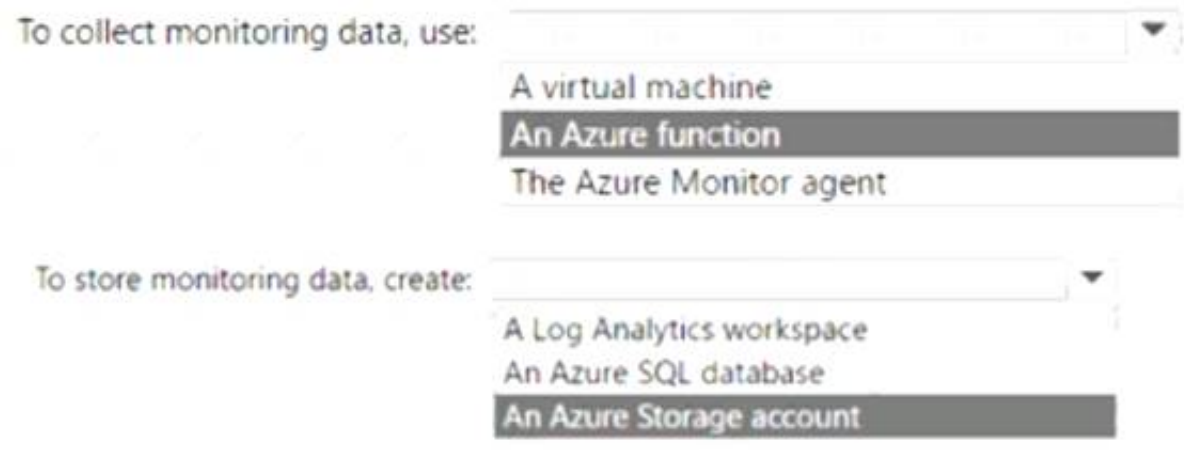

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1 = Azure Monitor Agent Box 2 = An Azure SQL database
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql

NEW QUESTION 10
- (Exam Topic 5)
You 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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

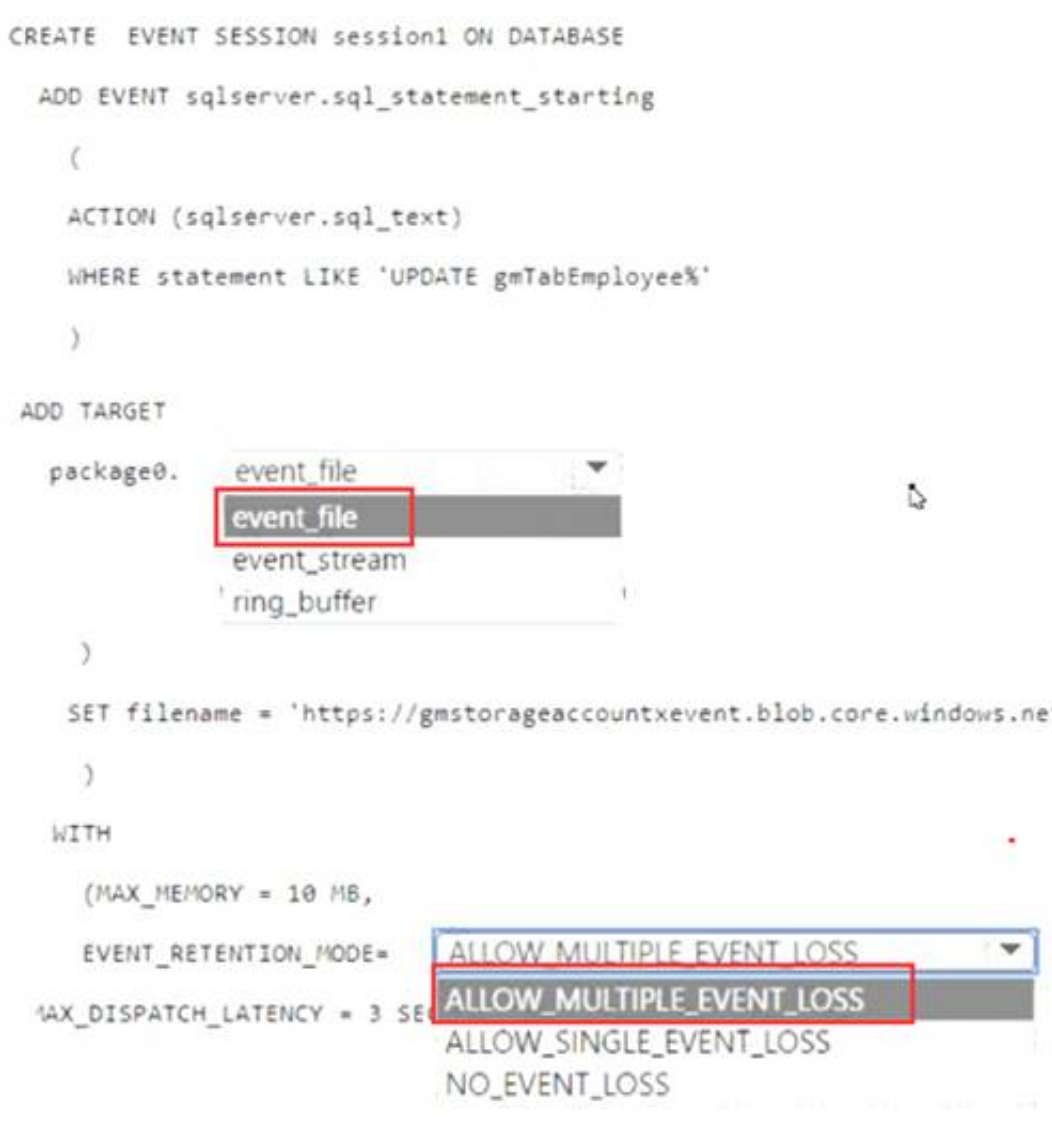

NEW QUESTION 13
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines named SQL1. SQL1 contains a database named DB1.
You need to enable Transparent Data Encryption (TDE) for DB1.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the

correct order.

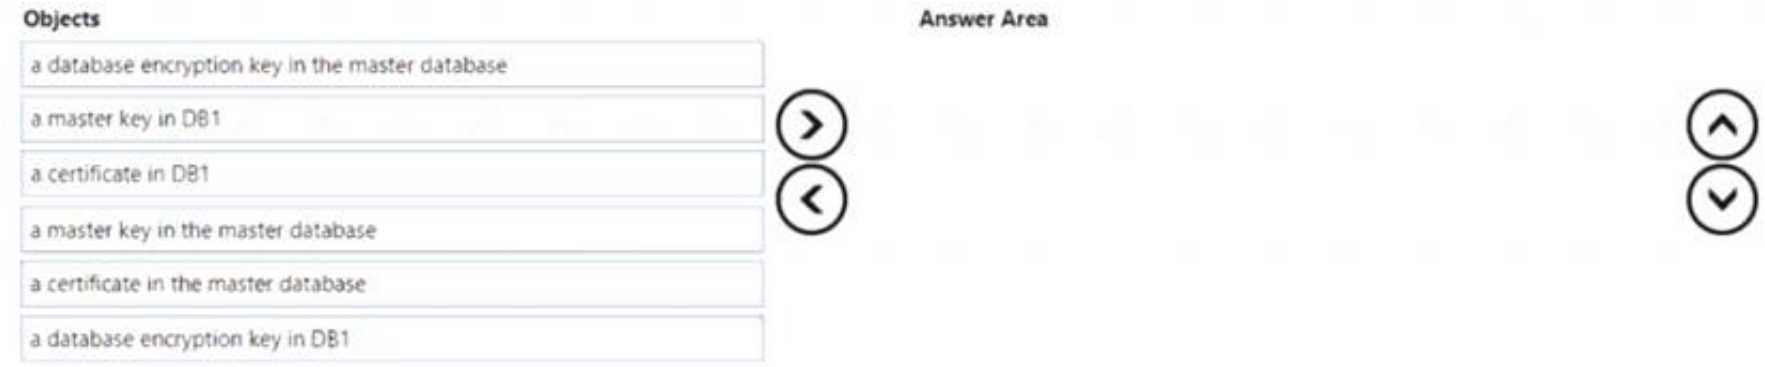

A. Mastered
B. Not Mastered

Answer: A

Explanation:

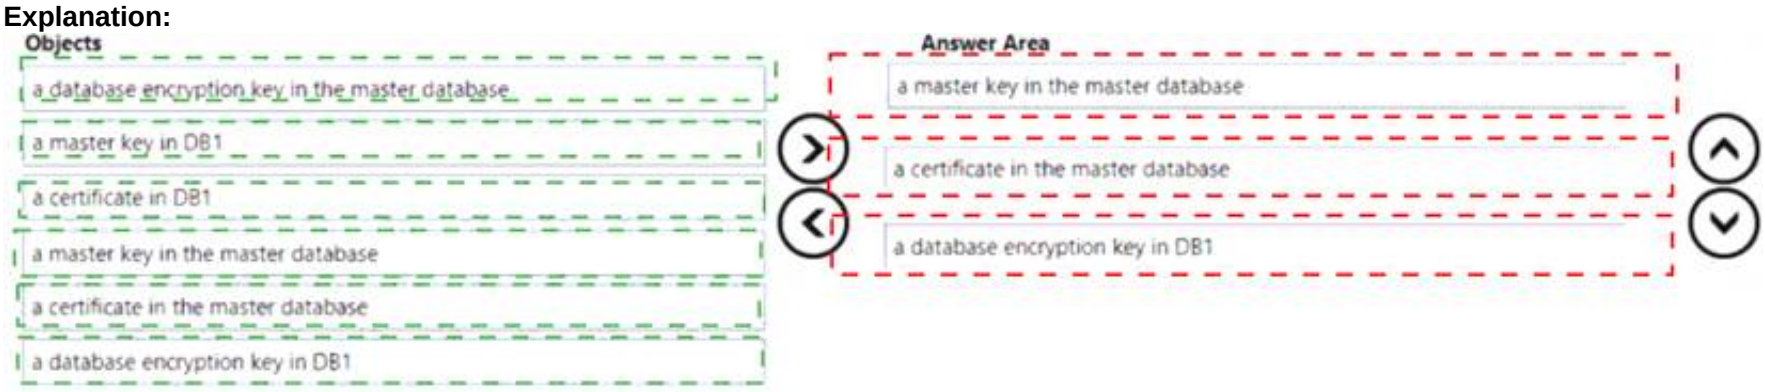

NEW QUESTION 14
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.
You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.
How should you configure the backup of DB1? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

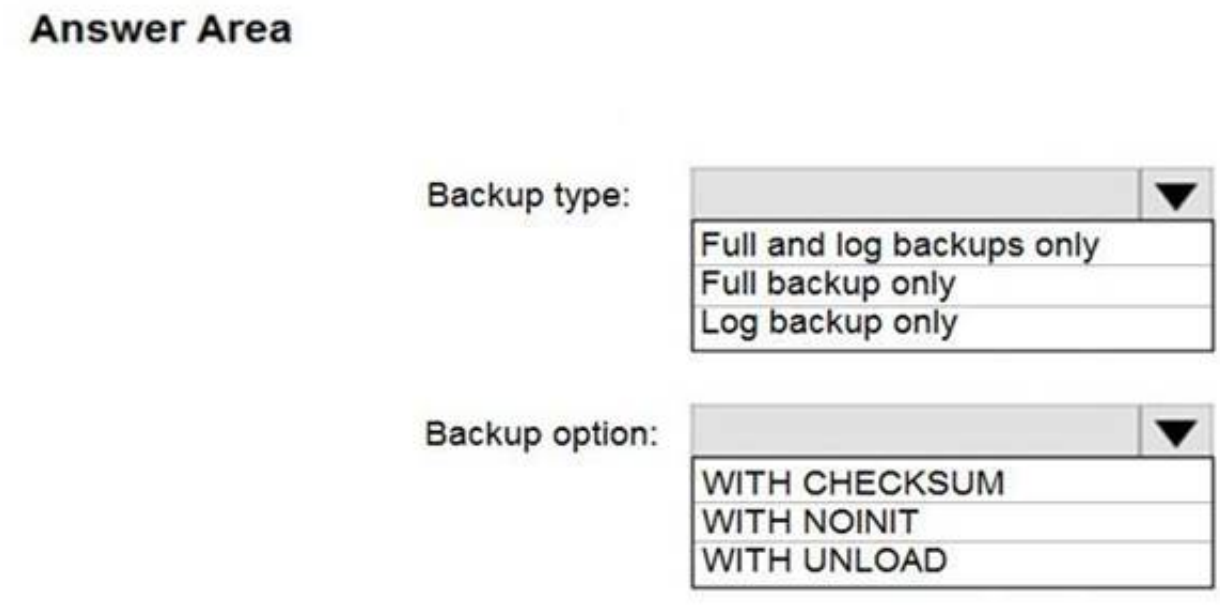

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Full and log backups only
Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.
Box 2: WITH CHECKSUM
Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.
Reference:
https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online

NEW QUESTION 19
- (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 21
- (Exam Topic 5)
You have an Azure Active Directory (Azure AD) tenant named contoso.com that contains a user named user1@contoso.com and an Azure SQL managed instance named SQLMI1.
You need to ensure that user1@contoso.com can create logins in SQLMI1 that map to Azure AD service principals.
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.

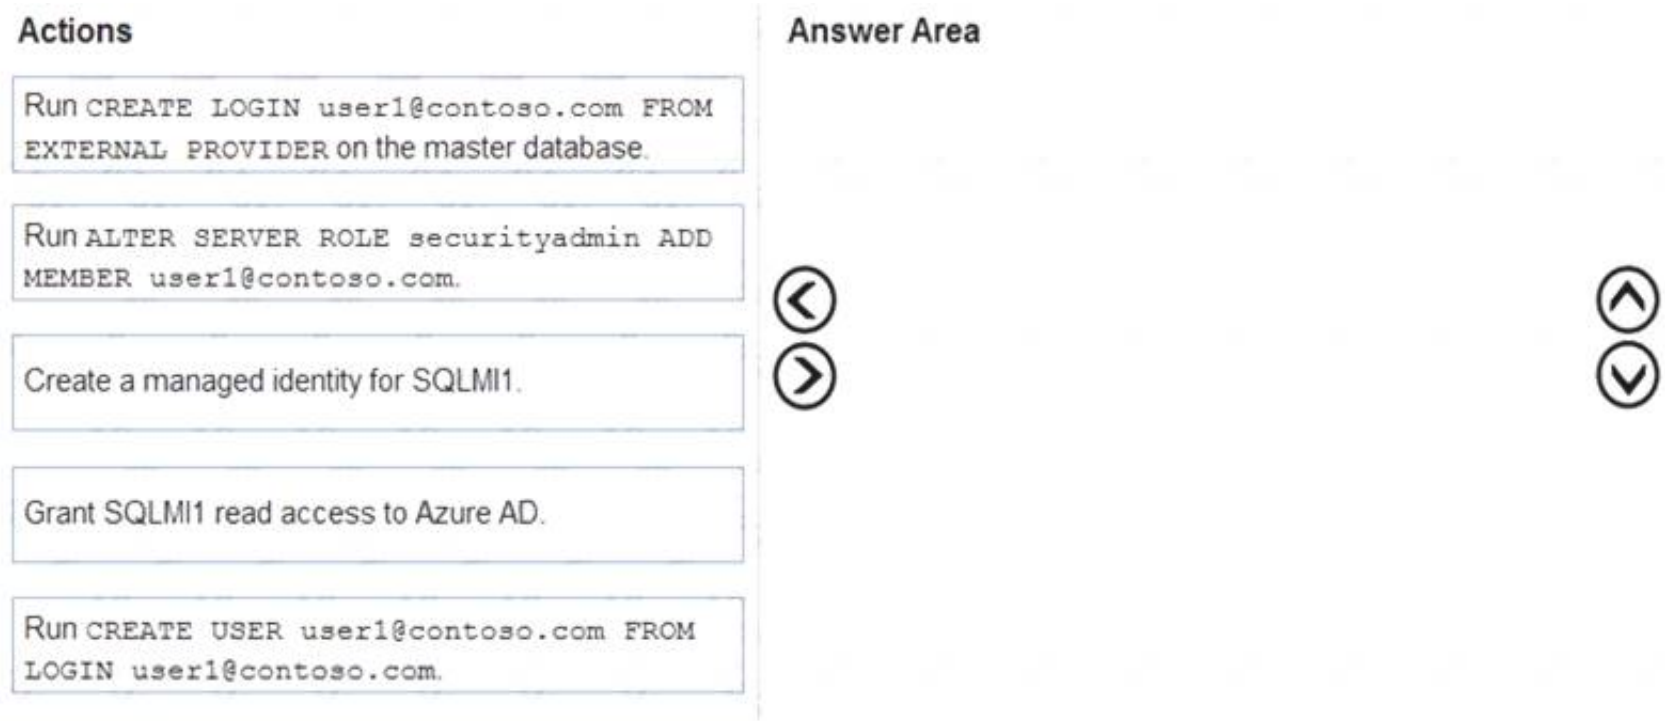

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated with medium confidence
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial

NEW QUESTION 26
- (Exam Topic 5)
You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:
- ProductID
- ItemPrice
- LineTotal
- Quantity
- StoreID
- Minute
- Month
- Hour
- Year
- Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: .partitionBy Example:
df.write.partitionBy("y","m","d") mode(SaveMode.Append) parquet("/data/hive/warehouse/db_name.db/" + tableName) Box 2:
("Year","Month","Day","Hour","StoreID")
Box 3: .parquet("/Purchases") Reference:
https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partiti

NEW QUESTION 30
- (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 31
- (Exam Topic 5)
You have an Azure Data Lake Storage Gen2 container.
Data is ingested into the container, and then transformed by a data integration application. The data isNOT modified after that. Users can read files in the container but cannot modify the files.
You need to design a data archiving solution that meets the following requirements:
- New data is accessed frequently and must be available as quickly as possible.
- Data that is older than five years is accessed infrequently but must be available within one second when requested.
- Data that us older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
- Costs must be minimized while maintaining the required availability.

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

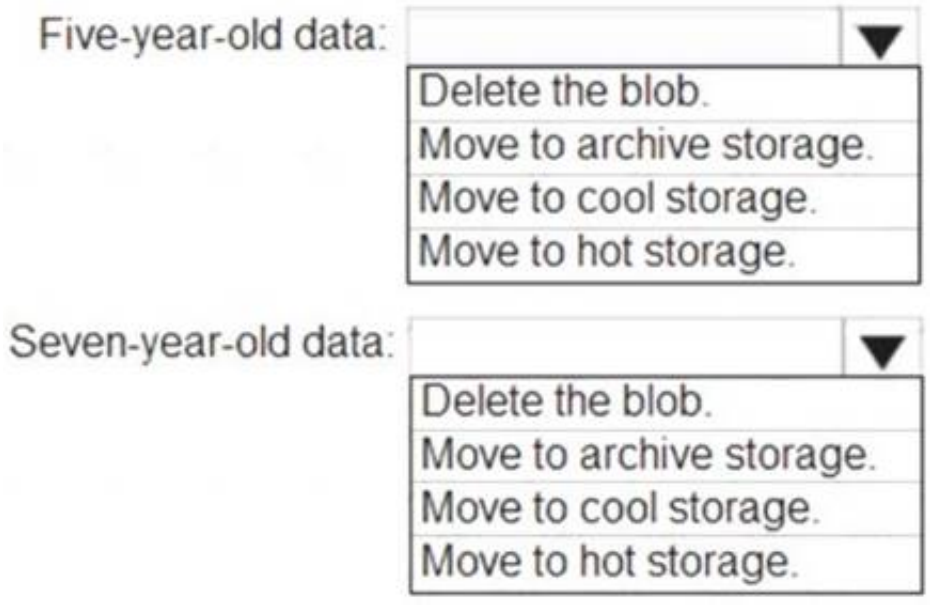

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text, table Description automatically generated
Box 1: Move to cool storage
The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:
Short-term backup and disaster recovery
Older data not used frequently but expected to be available immediately when accessed
Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot - Optimized for storing data that is accessed frequently.
Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.
Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.
Box 2: Move to archive storage
Example usage scenarios for the archive access tier include: Long-term backup, secondary backup, and archival datasets
Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers

NEW QUESTION 33
- (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 16 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: D

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 38
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1.
Which isolation level should you configure?

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

Answer: B

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

NEW QUESTION 42

- (Exam Topic 5)
You are monitoring an Azure Stream Analytics job.
You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero. You need to ensure that the job can handle all the events.
What should you do?

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

Answer: D

Explanation:
Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring

NEW QUESTION 44
- (Exam Topic 5)
You plan to migrate on-premises Microsoft SQL Server databases to Azure.
You need to identify which deployment and resiliency options meet the following requirements:
- Support user-initiated backups.
- Support multiple automatically replicated instances across Azure regions.
- Minimize administrative effort to implement and maintain business continuity. What should you identify? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

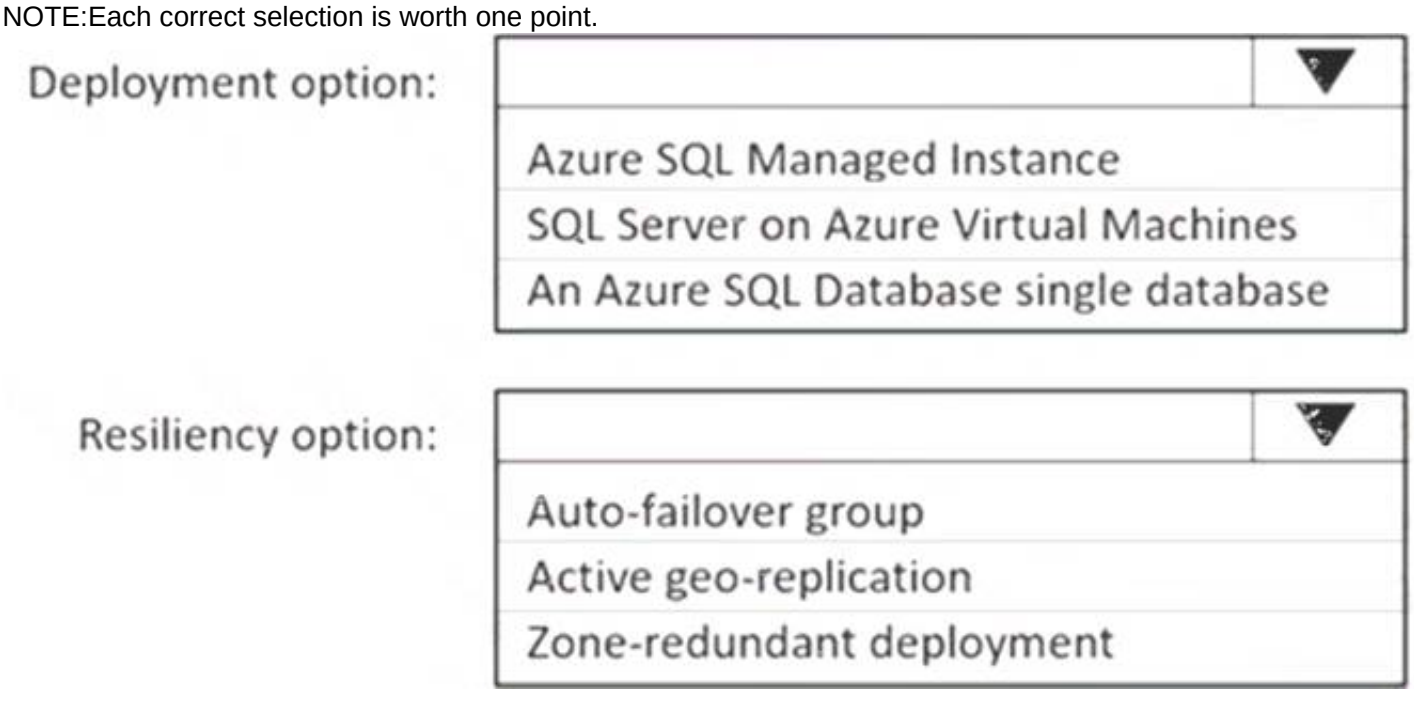

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SQL Server on Azure VMs
SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.
Box 2: Active geo-replication
Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-i https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure

NEW QUESTION 46
- (Exam Topic 5)
You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.
You need to ensure that an automatic email notification is sent once the job completes. What should you include in the solution?

A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
B. From SQL Server Management Studio (SSMS), runsp_set_sqlagent_properties
C. From SQL Server Management Studio (SSMS), create a Database Mail profile
D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

Answer: C

Explanation:
To send a notification in response to an alert, you must first configure SQL Server Agent to send mail. Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:
- In Object Explorer, expand a SQL Server instance.

- Right-click SQL Server Agent, and then click Properties.
- Click Alert System.
- Select Enable Mail Profile.
- In the Mail system list, select Database Mail.
- In the Mail profile list, select a mail profile for Database Mail.
- Restart SQL Server Agent.

Note: Prerequisites include:

- Enable Database Mail.
- Create a Database Mail account for the SQL Server Agent service account to use.
- Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.
- Set the profile as the default profile for the msdb database. Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d

NEW QUESTION 49
- (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 50
- (Exam Topic 5)
You have an Azure SQL Database managed instance. The instance starts experiencing performance issues.
You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.
What should you use?

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

Answer: D

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

NEW QUESTION 54
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change.
You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.
The database fails on Wednesday at 10:00.
Which three backups should you restore in sequence? To answer, move the appropriate backups from the list of backups to the answer area and arrange them in the correct order.

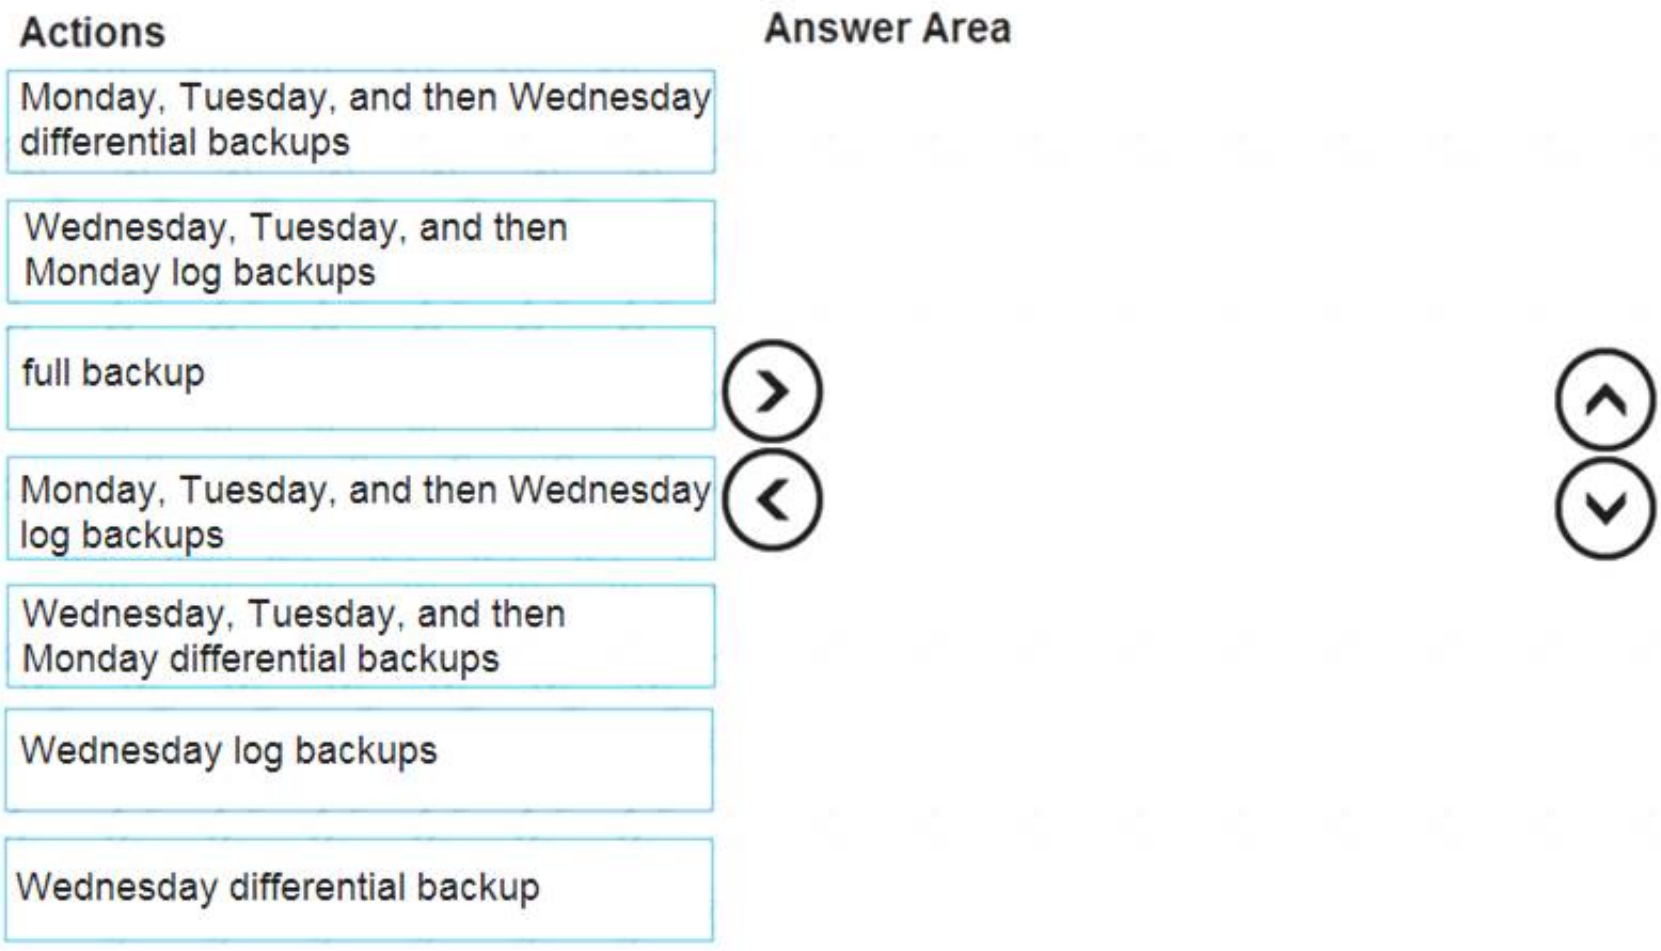

A. Mastered
B. Not Mastered

Answer: A

Explanation:

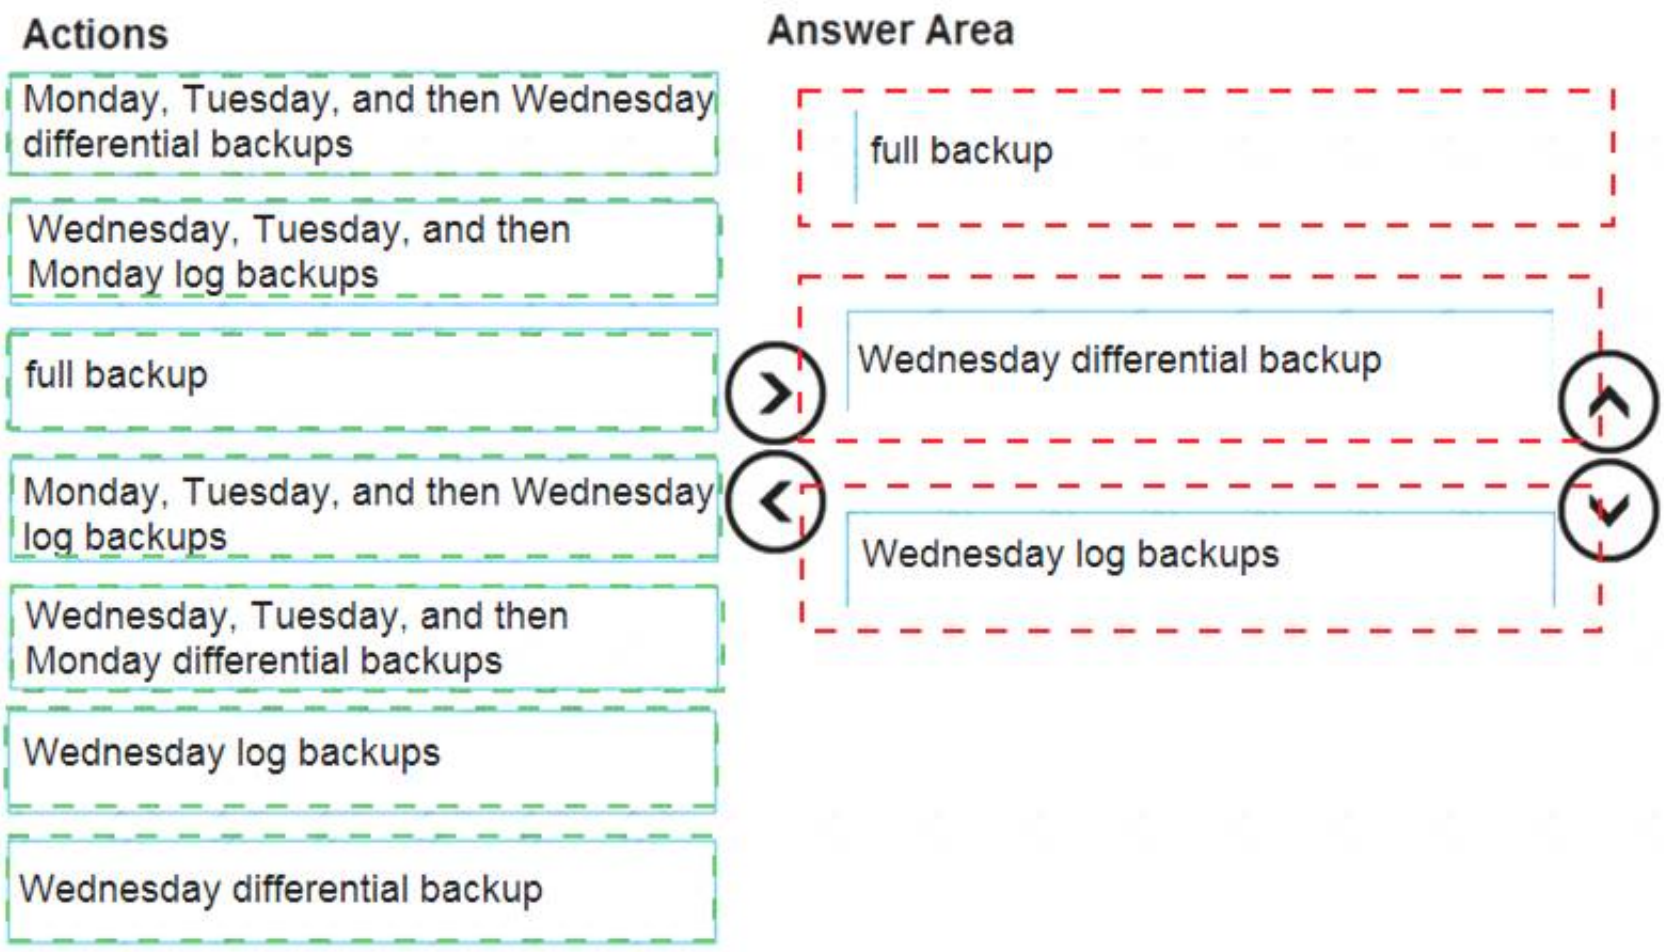

NEW QUESTION 58
- (Exam Topic 5)
You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

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

The records contain two applicants at most.
You need to build a table that includes only the address fields.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: CREATE EXTERNAL TABLE
An external table points to data located in Hadoop, Azure Storage blob, or Azure Data Lake Storage. External tables are used to read data from files or write data to files in Azure Storage. With Synapse SQL, you can use external tables to read external data using dedicated SQL pool or serverless SQL pool.
Syntax:
CREATE EXTERNAL TABLE { database_name.schema_name.table_name | schema_name.table_name | table_name }
(<column_definition> [,...n]) WITH (
LOCATION = 'folder_or_filepath', DATA_SOURCE = external_data_source_name, FILE_FORMAT = external_file_format_name Box 2. OPENROWSET
When using serverless SQL pool, CETAS is used to create an external table and export query results to Azure Storage Blob or Azure Data Lake Storage Gen2.
Example: AS
SELECT decennialTime, stateName, SUM(population) AS population FROM
OPENROWSET(BULK
'https://azureopendatastorage.blob.core.windows.net/censusdatacontainer/release/us_population_county/year=*/
FORMAT='PARQUET') AS [r]
GROUP BY decennialTime, stateName GO
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

NEW QUESTION 59
- (Exam Topic 5)
You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

A. automated
B. interactive
C. High Concurrency

Answer: A

Explanation:
Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.
The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.
Reference:
https://docs.microsoft.com/en-us/azure/databricks/clusters

NEW QUESTION 63
- (Exam Topic 5)
You have an Azure SQL database named db1.
You need to retrieve the resource usage of db1 from the last week.
How should you complete the statement? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

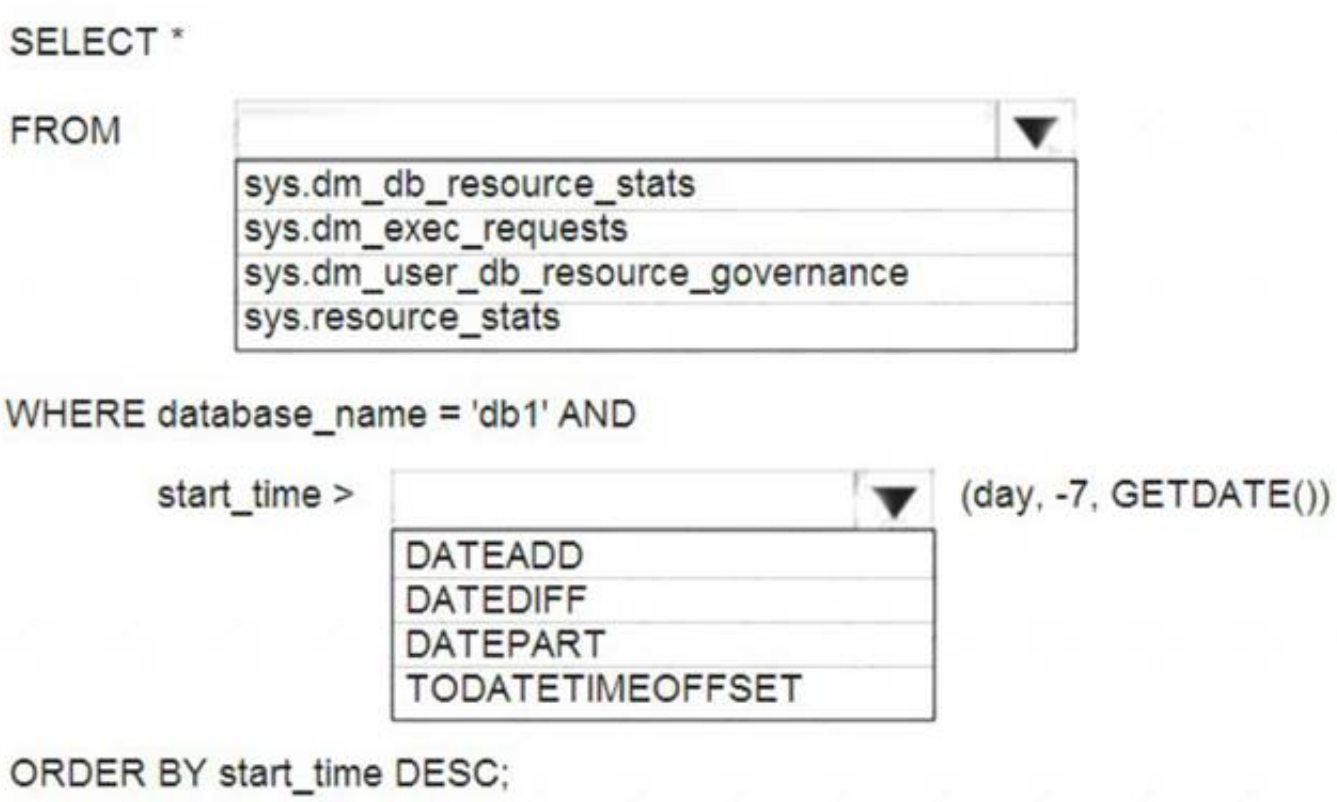

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: sys.resource_stats
sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.
Box 2: DateAdd
The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.
DECLARE @s datetime; DECLARE @e datetime;
SET @s= DateAdd(d,-7,GetUTCDate()); SET @e= GETUTCDATE();
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats
WHERE start_time BETWEEN @s AND @e GROUP BY database_name
HAVING AVG(avg_cpu_percent) >= 80 Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-data

NEW QUESTION 64
- (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.

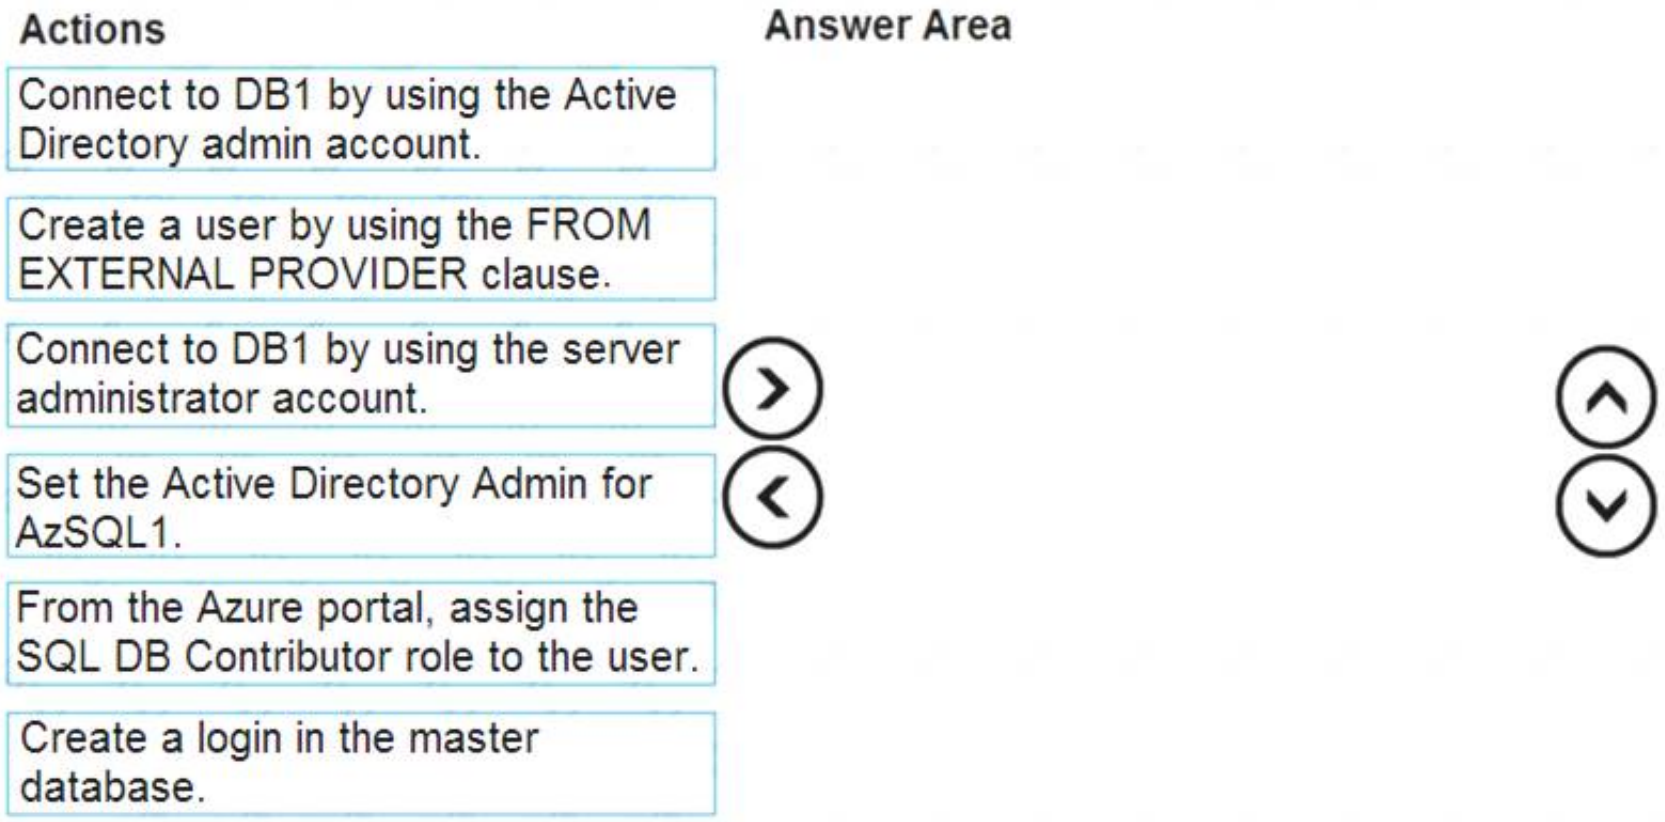

A. Mastered
B. Not Mastered

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 69
- (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.

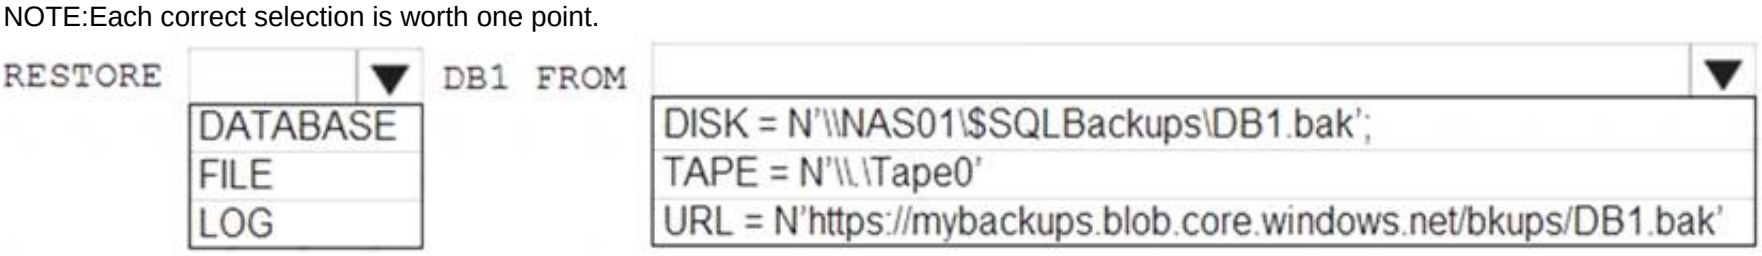

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated

NEW QUESTION 72
- (Exam Topic 5)
You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:
- The first six numerals of each customer's phone number must be masked.
- The last four digits of each customer's phone number must be visible.
- Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

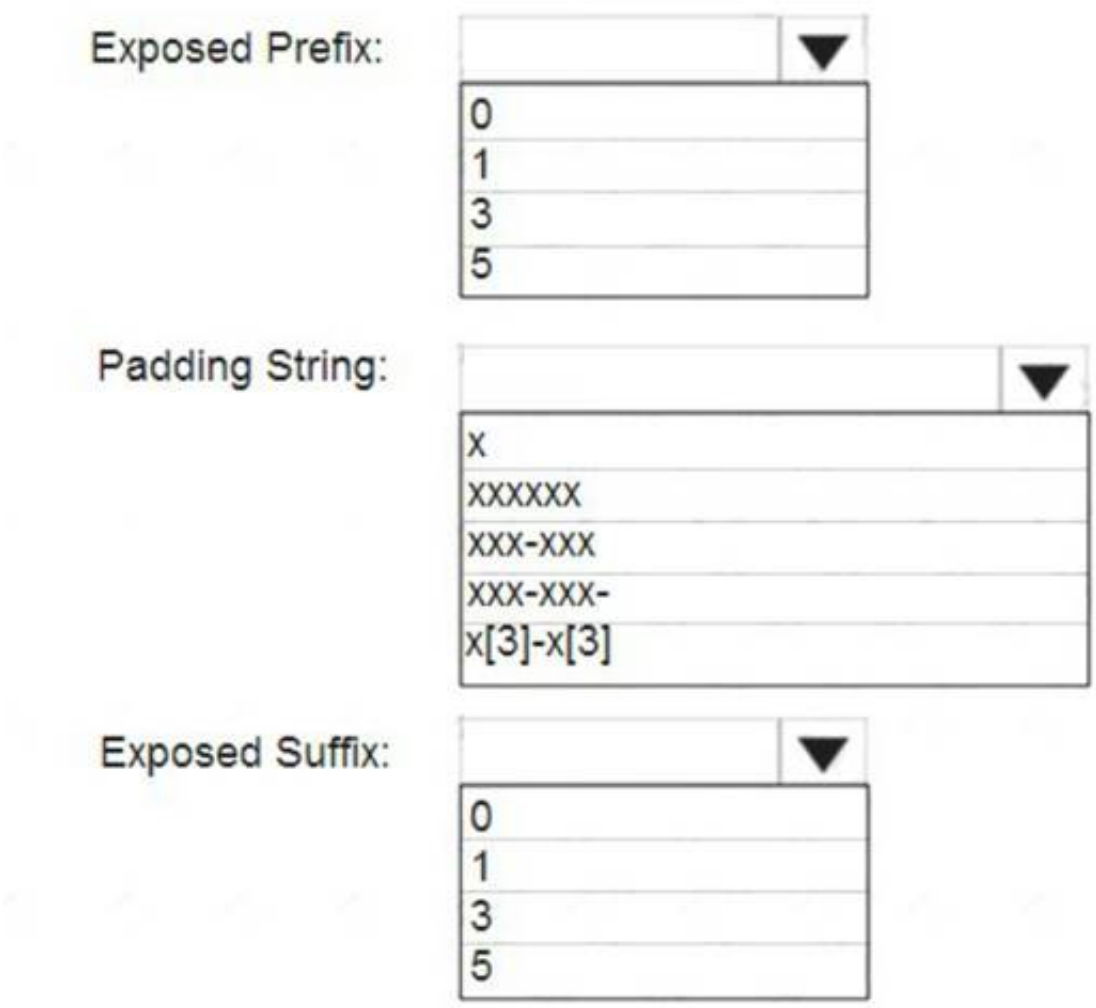

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: 0
Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix
Box 2: xxx-xxx
Box 3: 5 Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

NEW QUESTION 77
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine.
You need to add a 4-TB volume that meets the following requirements:

- Maximizes IOPs
- Uses premium solid state drives (SSDs)

What should you do?

A. Attach two mirrored 4-TB SSDs.
B. Attach a stripe set that contains four 1-TB SSDs.
C. Attach a RAID-5 array that contains five 1-TB SSDs.
D. Attach a single 4-TB SSD.

Answer: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=window

NEW QUESTION 81
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL managed instance named SQLMi1 and a SQL Agent job named Backupdb. Backupdb performs a daily backup of the databases hosted on SQLMi1.
You need to be notified by email if the job fails.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE:More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance

NEW QUESTION 82
- (Exam Topic 5)
You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL
You need to use PowerShell to enable and configure automated patching for SQL The solution must include both SQL Server and Windows security updates.
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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

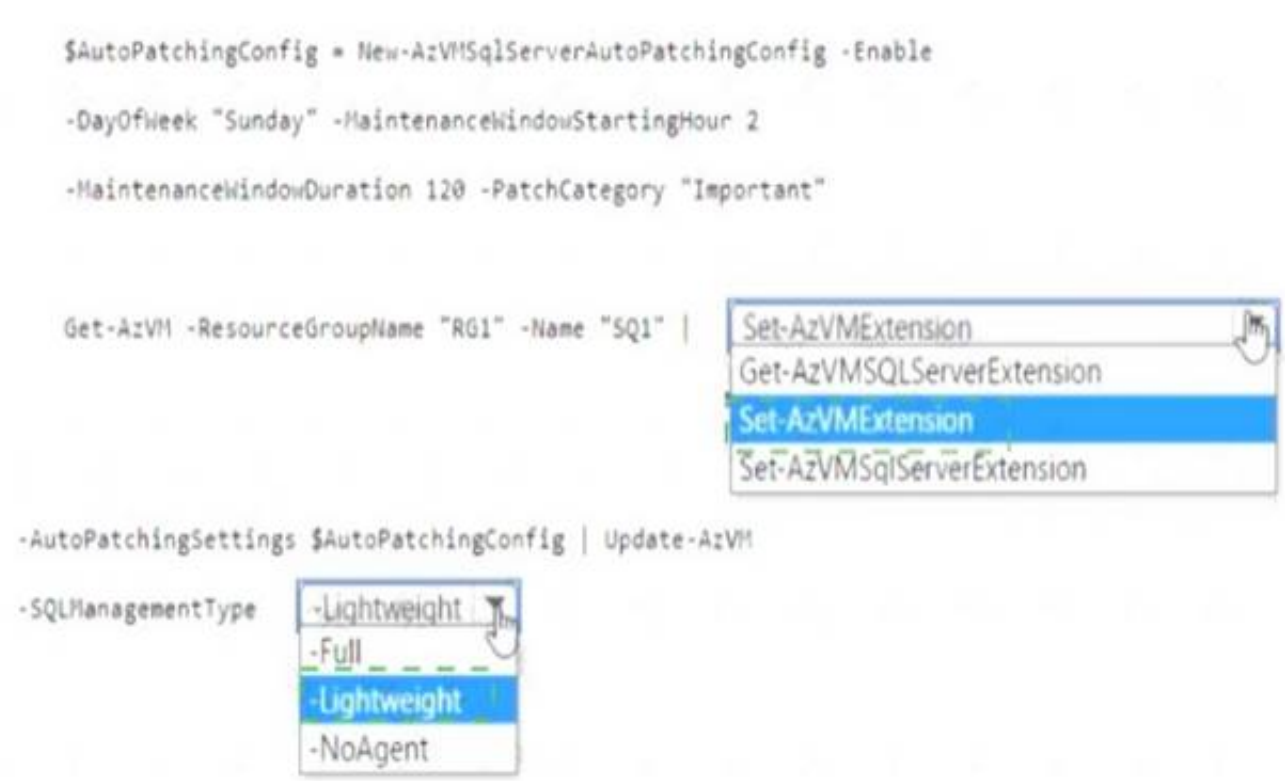

NEW QUESTION 86
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines.
You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The solution must use the principle of least privilege.
Which role should you assign to User1?

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

Answer: A

NEW QUESTION 90
- (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.

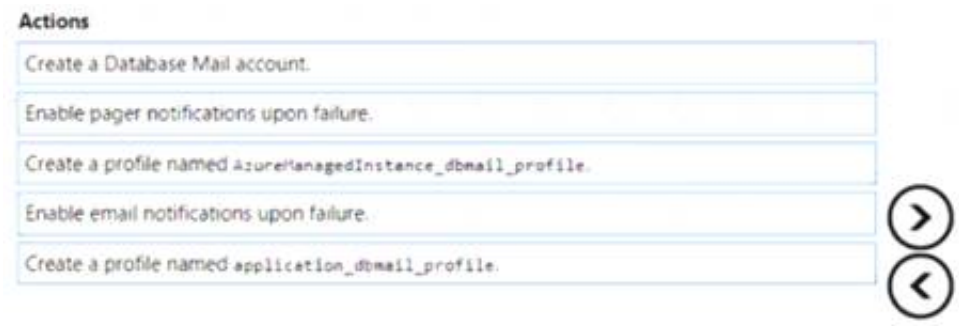

A. Mastered
B. Not Mastered

Answer: A

Explanation:

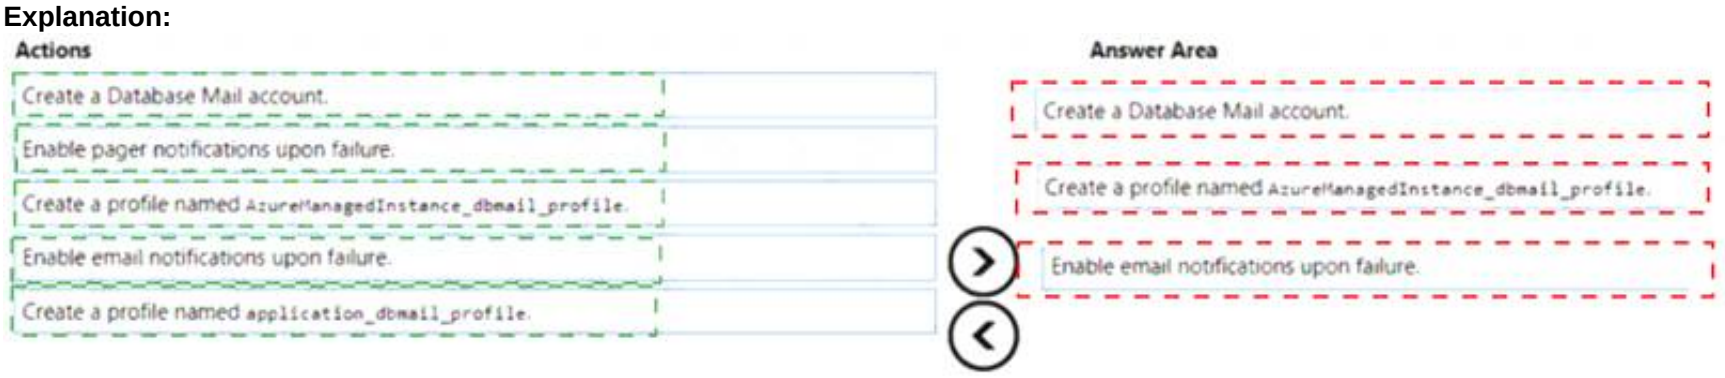

NEW QUESTION 92
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named Db1. You need to enable automatic tuning for Db1.
How should you complete the statements? To answer, select the appropriate answer in the answer area.
NOTE:Each correct selection is worth one point.

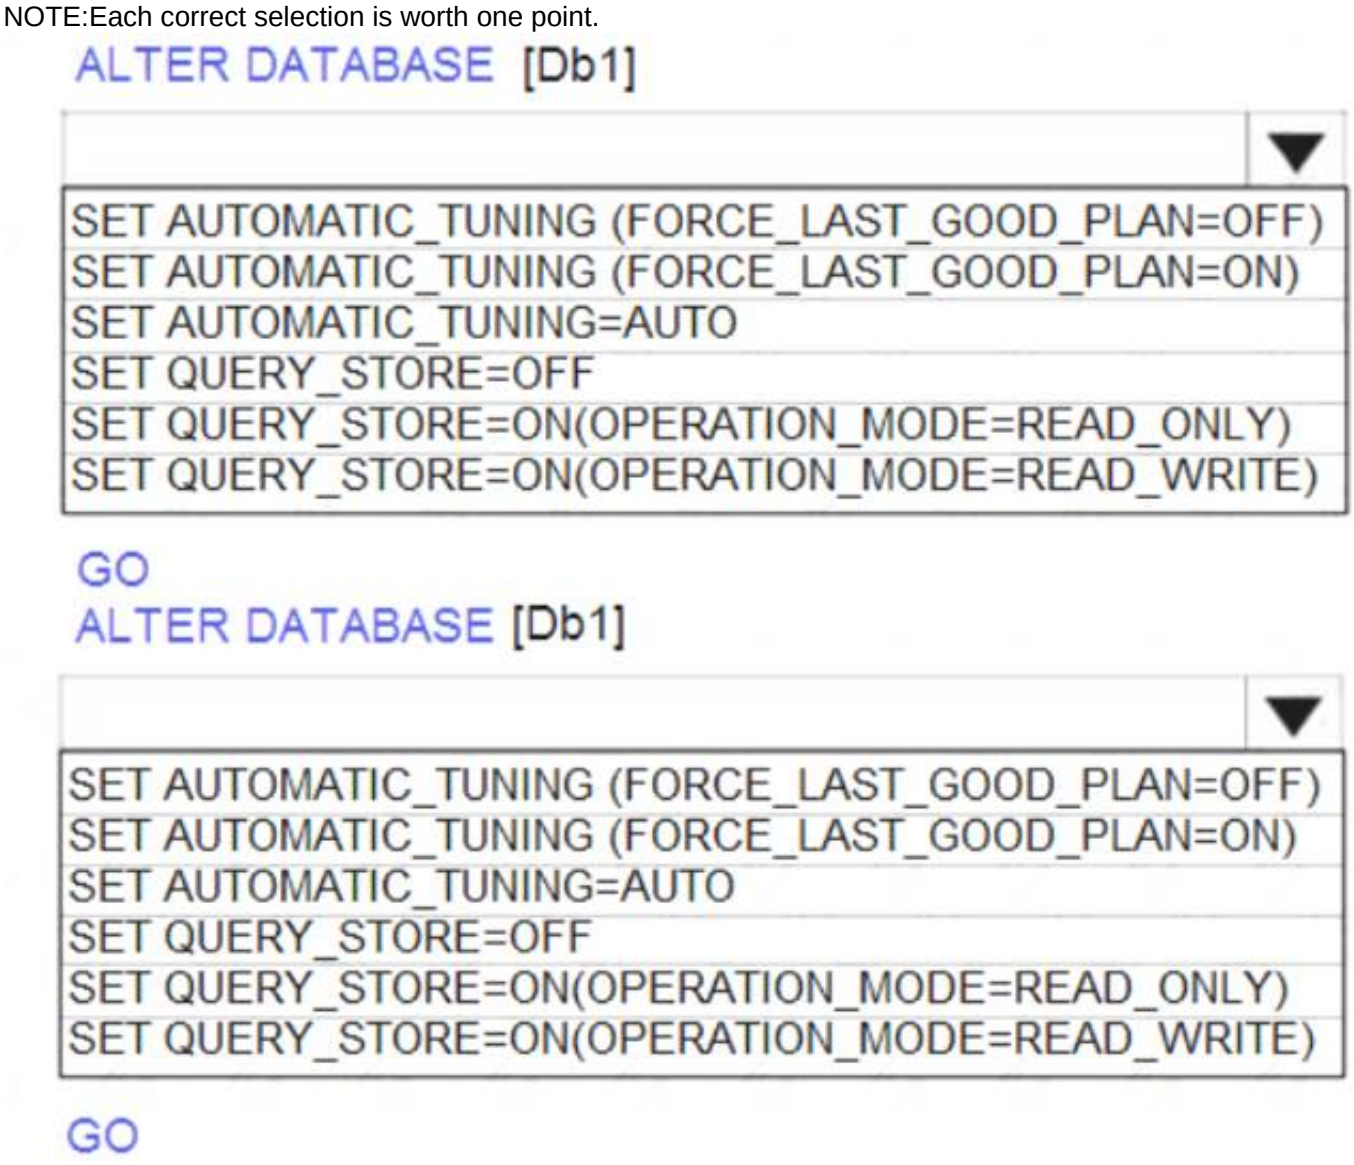

A. Mastered
B. Not Mastered

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 94
- (Exam Topic 5)
You have an Azure subscription that contains a SQL Server on Azure Virtual Machines instance named SQLVMI. SQLVMI hosts a database named OBI.
You need to retrieve query plans from the Query Store on DBI. What should you do first?

A. On SQLVM1, install the SQL Server laaS Agent extension.
B. From Microsoft SQL Server Management Studio, modify the properties of the SQL Server instance.
C. From Microsoft SQL Server Management Studio, modify the properties of DB 1.
D. On SQLVM1, install the Azure Monitor agent for Windows.

Answer: B

NEW QUESTION 97
- (Exam Topic 5)
You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8. You need to deploy a failover cluster instance (FCI) to VM1 and VM2. The solution must eliminate the need for the following:
• A distributed network name (DNN)
• A load balancer What should you do?

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

Answer: D

NEW QUESTION 101
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.
You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.
Solution: You restore Database1 from Server1 to the Server2 by using theRESTORETransact-SQL command and theREPLACEoption.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
The REPLACE option overrides several important safety checks that restore normally performs. The overridden checks are as follows:
- Restoring over an existing database with a backup taken of another database.

With the REPLACE option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

NEW QUESTION 102
- (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 107
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.
You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed. You need to ensure that webapp1 can identity the changes to Table'. The solution must meet the following
requirements:
• Minimize compute times.
• Minimize storage.
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.

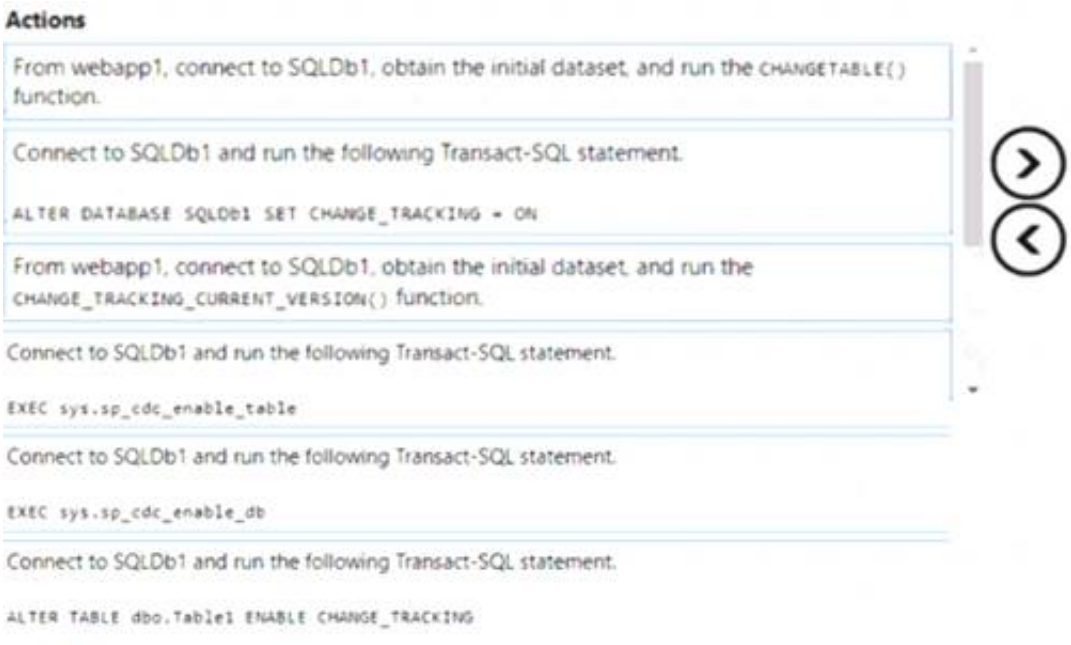

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated

NEW QUESTION 112
- (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 113
- (Exam Topic 5)
You have an Azure SQL Database instance named DatabaseA on a server named Server1.
You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.
You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

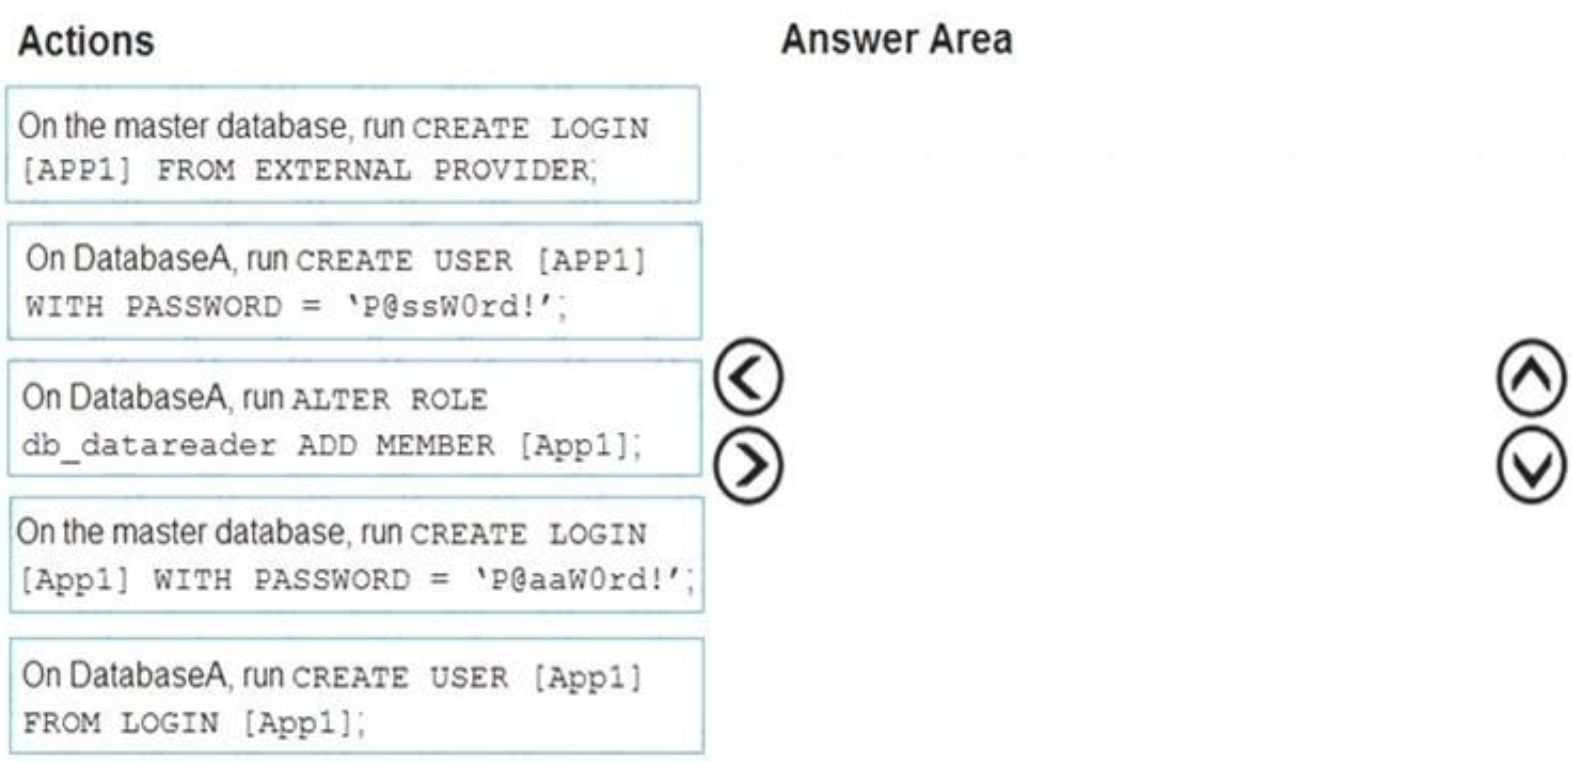

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.
Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]
Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:
CREATE USER readonlyuser FROM LOGIN readonlylogin;
Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]
Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the db_datareader role.
EXEC sp_addrolemember 'db_datareader', 'readonlyuser'; Reference:
https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/

NEW QUESTION 114
- (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 115
- (Exam Topic 5)
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

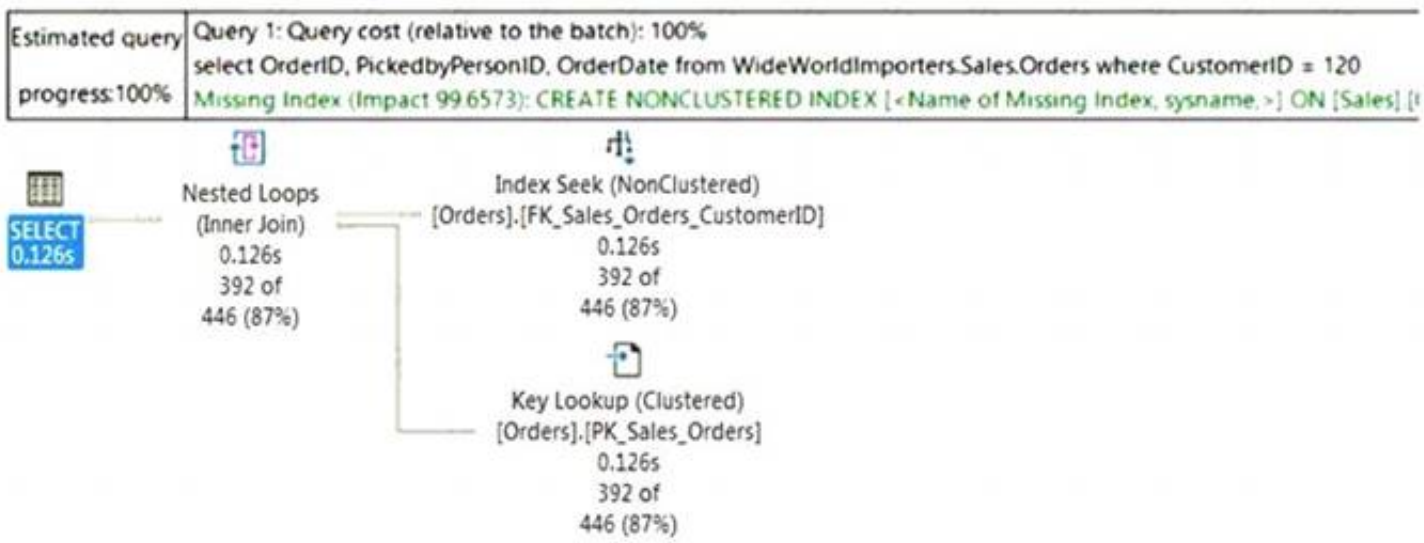

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE:Each correct selection is worth one point.

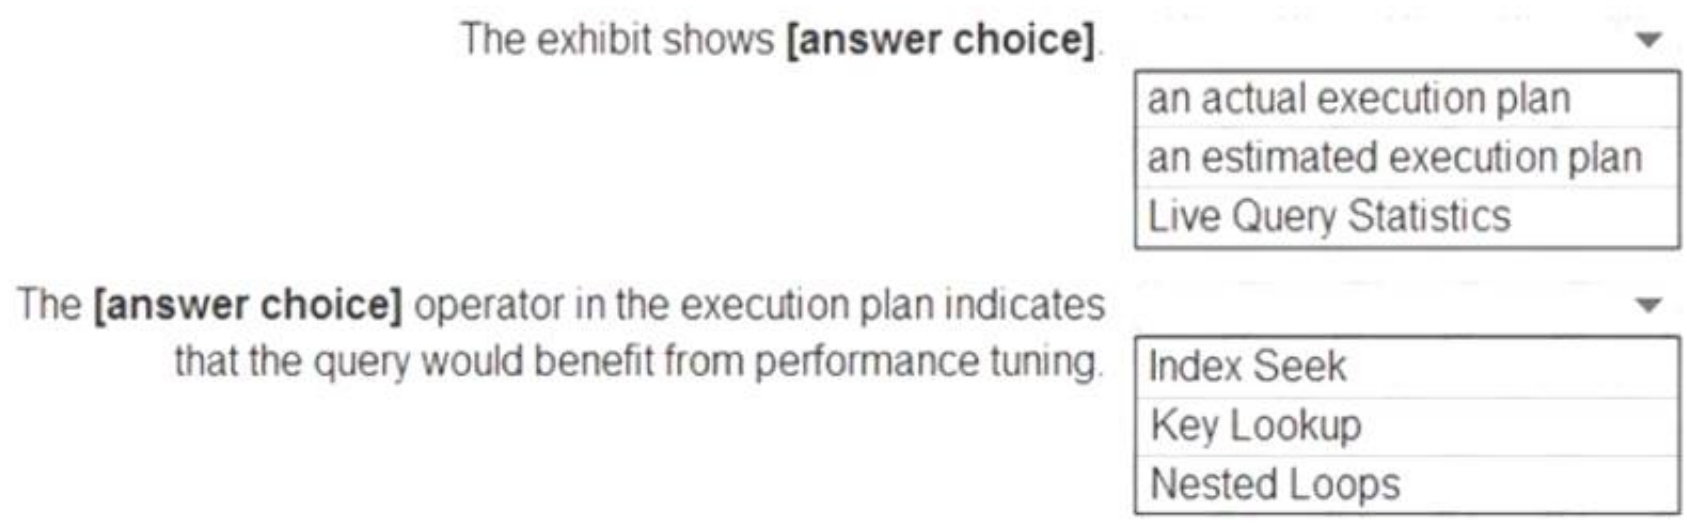

A. Mastered
B. Not Mastered

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 119
- (Exam Topic 5)
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.
You plan to deploy an Always On availability group that will have the following configurations:
- VM1 will host the primary replica of DB1.
- VM2 will host a secondary replica of DB1.

You need to prepare the secondary database on VM2 for the availability group.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

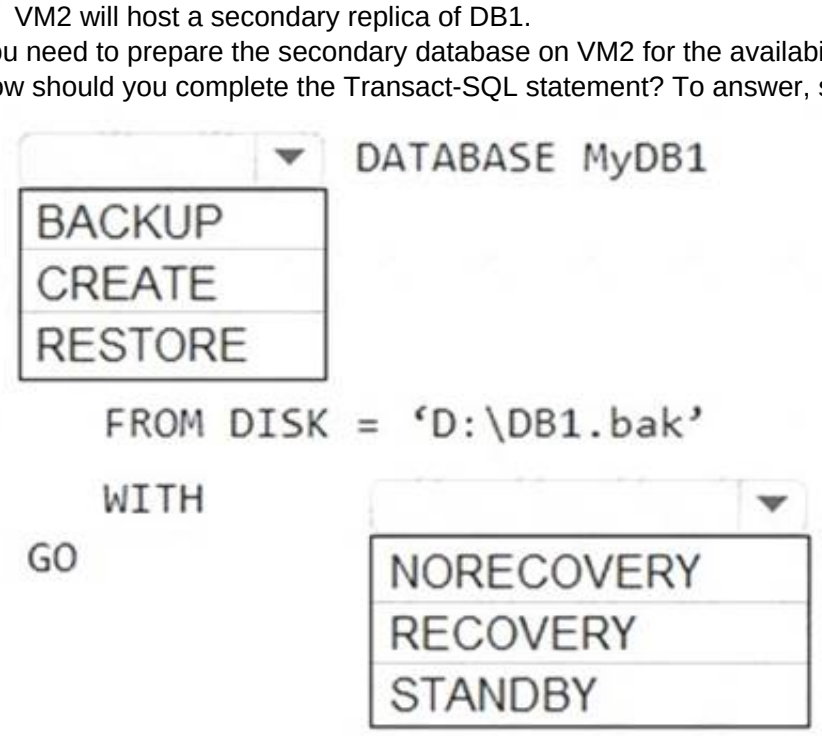

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, chat or text message Description automatically generated
Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondar

NEW QUESTION 122
- (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 127
- (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.

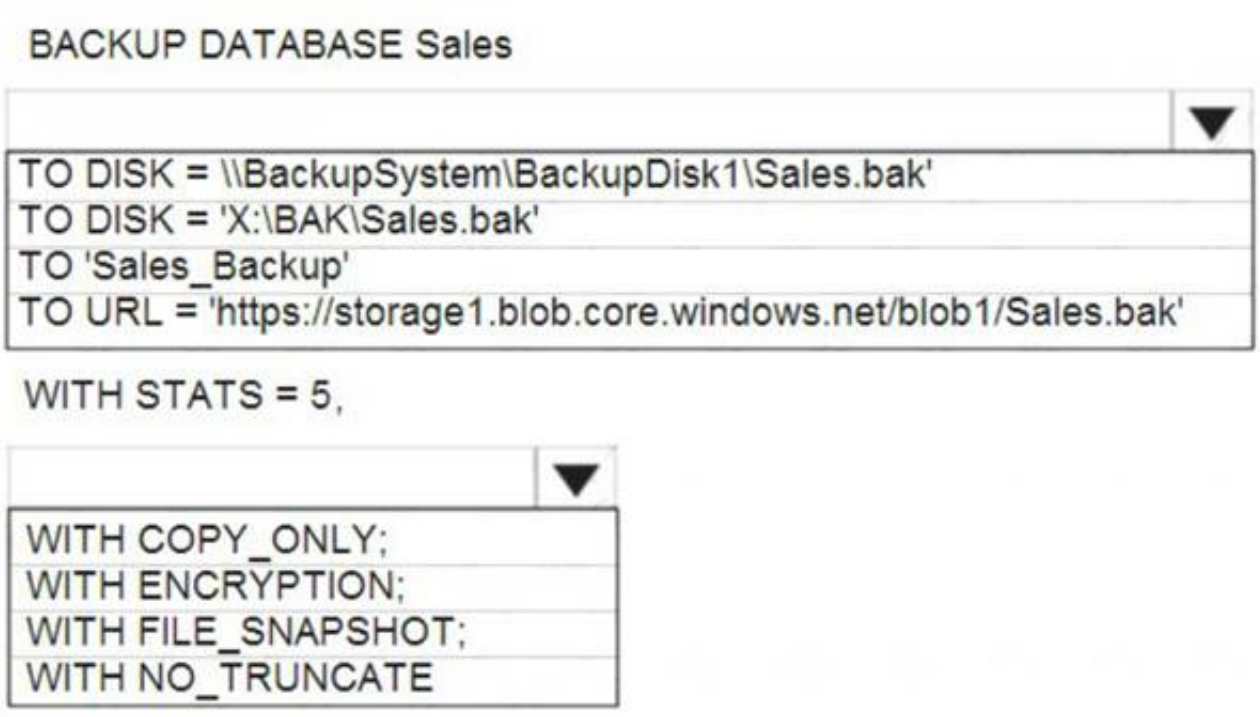

A. Mastered
B. Not Mastered

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 128
- (Exam Topic 5)
You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.
You need to ensure that the design provides the fastest query times of the records when querying for arbitrary date ranges and aggregating by fiscal calendar attributes.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Answer: BD

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

NEW QUESTION 132
- (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 137
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy a logical SQL server by using PowerShell. The solution must ensure that the logical SQL server can create Azure AD users and provide Transparent Data Encryption (TDE} with a customer-managed key.
How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

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

- assignidentity
-federatedclientID
-keyid

- assignidentity
-federatedclientID
-keyid

A. Mastered
B. Not Mastered

Answer: A

Explanation:
"Assigned Identity" and "Key
id"https://learn.microsoft.com/en-us/powershell/module/az.sql/new-azsqlserver?view=azps-10.2.0#code-try-3

NEW QUESTION 140
- (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 144
- (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 148
- (Exam Topic 5)
You have an Azure subscription.
Vou need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:
• Minimize how long it takes to fail over.
• Maintain existing connections to the primary replica during a failover. What should you do?

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

Answer: C

NEW QUESTION 150
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

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

You need to configure a connection between VM1 and MIL The solution must meet the following requirements:
• The connection must be encrypted.
• Network latency must be minimized. What should you implement?

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

Answer: B

NEW QUESTION 154
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image.
You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.
What should you do first?

A. From the SQL Server Installation Center on SQLVM1. run the Edition Upgrade wizard.
B. From SQLVM1, uninstall the SQL Server instance.
C. From the SQL Server Installation Center on SQLVM1. run the Repair wizard.
D. From the Azure portal, reconfigure SQLVM1.

Answer: B

NEW QUESTION 157
- (Exam Topic 5)
You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1. You need to ensure that User1 can create databases on Server1. Which database role should you assign to User1?

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

Answer: B

NEW QUESTION 158
- (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 160
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance named VM1 . You run the following query.

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

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

Answer Area

Statements	Yes	No
The log file will be truncated.	○	○
DB1 will be placed in an offline state.	○	○
You are performing a tail-log backup.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The log file will be truncated.	○	○
DB1 will be placed in an offline state.	○	○
You are performing a tail-log backup.	○	○

NEW QUESTION 164
- (Exam Topic 5)
You configure backups for an Azure SQL database as shown in the following exhibit.

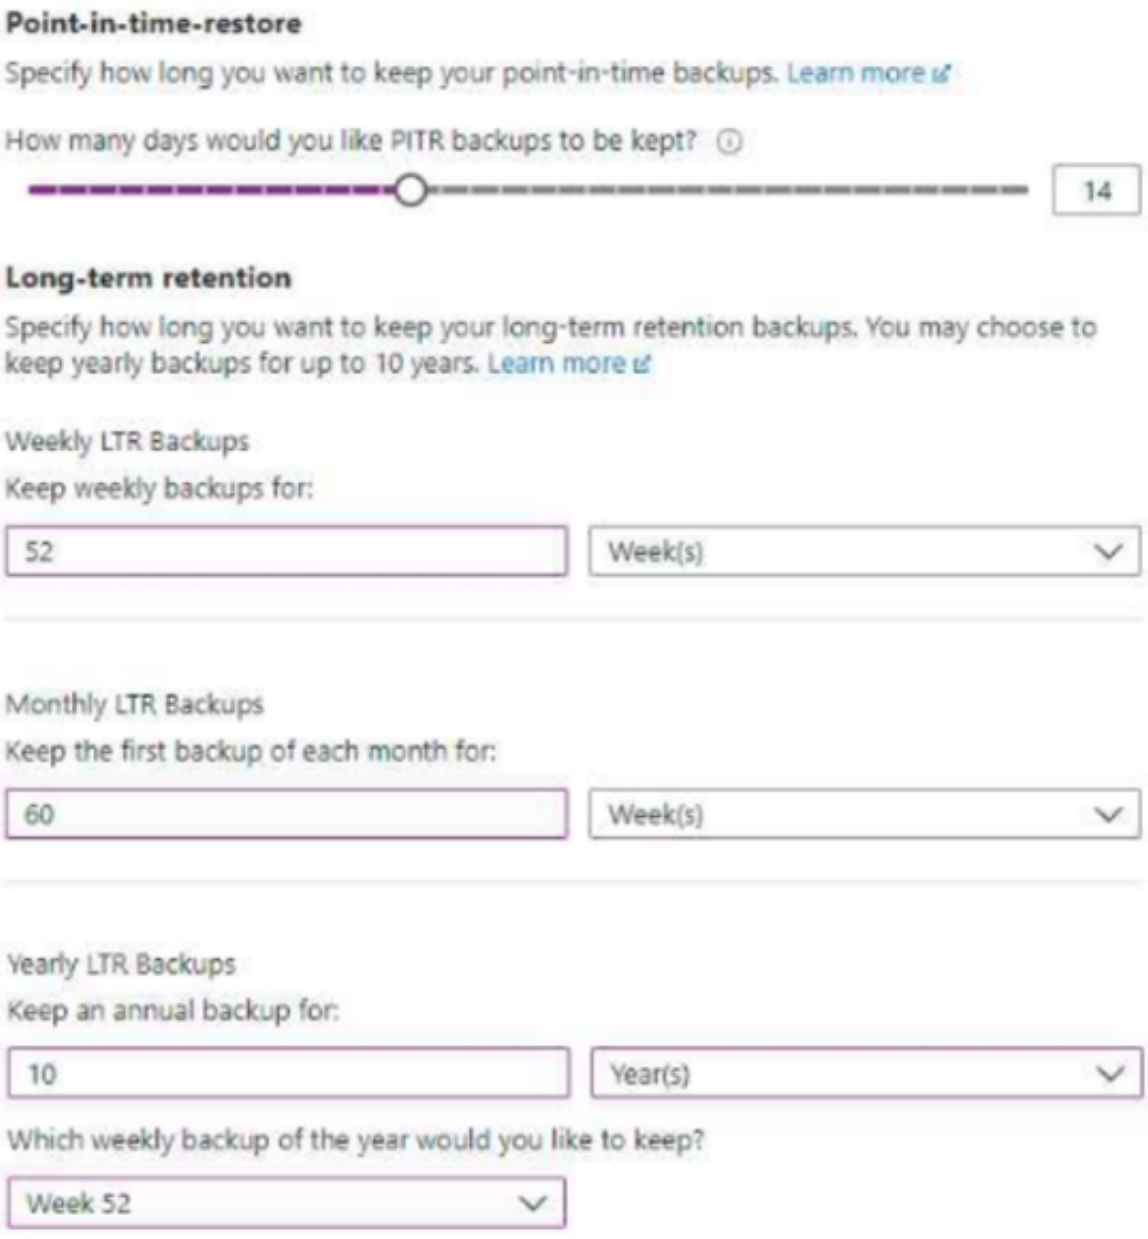

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a **[answer choice]**
- point-time restore (PITR) backup.
- yearly long-term retention (LTR) backup.
- weekly long-term retention (LTR) backup.
- monthly long-term retention (LTR) backup.

After the 52nd weekly backup runs, there will be **[answer choice]** in long term retention.
- 1 backup copy
- 52 backup copies
- 64 backup copies
- 65 backup copies

NEW QUESTION 166
- (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 168
- (Exam Topic 5)
You manage an enterprise data warehouse in Azure Synapse Analytics.
Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.
You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

A. Local tempdb percentage
B. DWU percentage
C. Data Warehouse Units (DWU) used
D. Cache hit percentage

Answer: A

Explanation:
Tempdb is used to hold intermediate results during query execution. High utilization of the tempdb database can lead to slow query performance.
Note: If you have a query that is consuming a large amount of memory or have received an error message related to allocation of tempdb, it could be due to a very large CREATE TABLE AS SELECT (CTAS) or INSERT SELECT statement running that is failing in the final data movement operation.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-managemonit

NEW QUESTION 172
- (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.

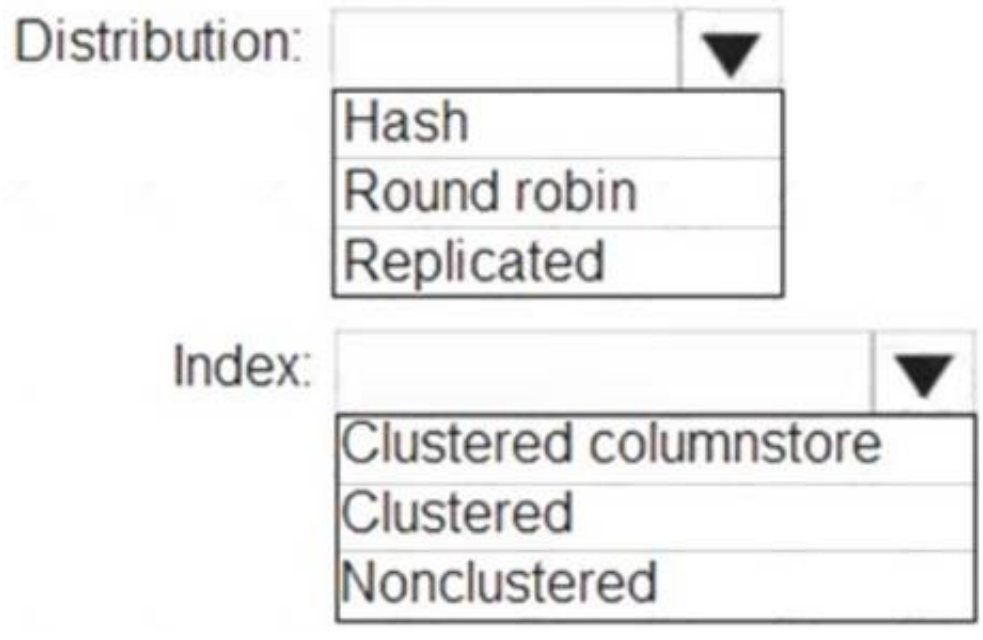

A. Mastered
B. Not Mastered

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 177
- (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 179
- (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 182
- (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 187
- (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 190
- (Exam Topic 5)
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?

A. two Azure SQL Databases in an elastic pool
B. two databases hosted in SQL Server on an Azure virtual machine
C. two databases in an Azure SQL Managed instance

D. two single Azure SQL databases

Answer: A

Explanation:
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

NEW QUESTION 193
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes mapping data flow, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not a mapping flow,5 with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed. Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

NEW QUESTION 196
- (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 200
- (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_addremoteIogin
B. xp_instance_regwrite
C. sp_cnarge_users_login
D. xp_grant_login

Answer: B

NEW QUESTION 205
- (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 210
- (Exam Topic 5)
You have an Azure SQL database.
You have a query and the associated execution plan as shown in the following exhibit.

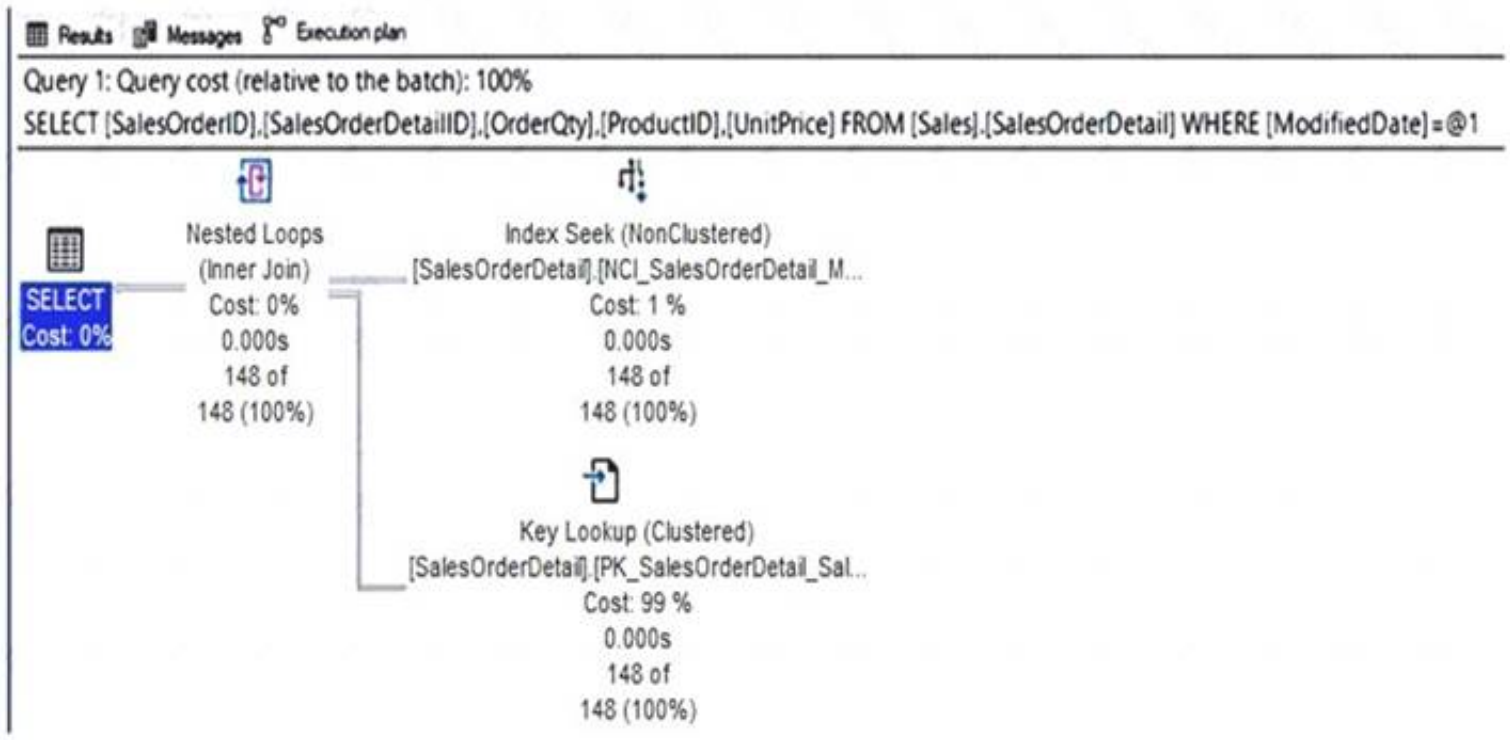

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE:Each correct selection is worth one point.

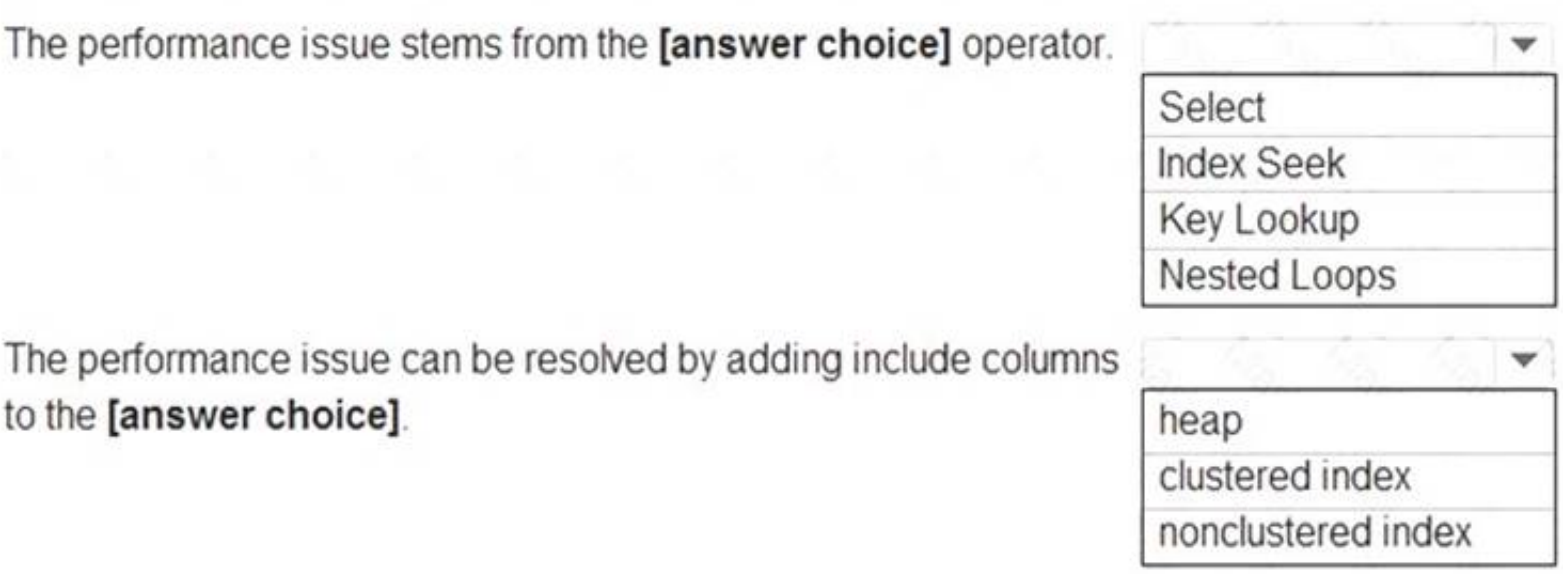

A. Mastered
B. Not Mastered

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 215
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil.
You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.
You need to create the backups for the migration. The solution must minimize the number of backup files created.
Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

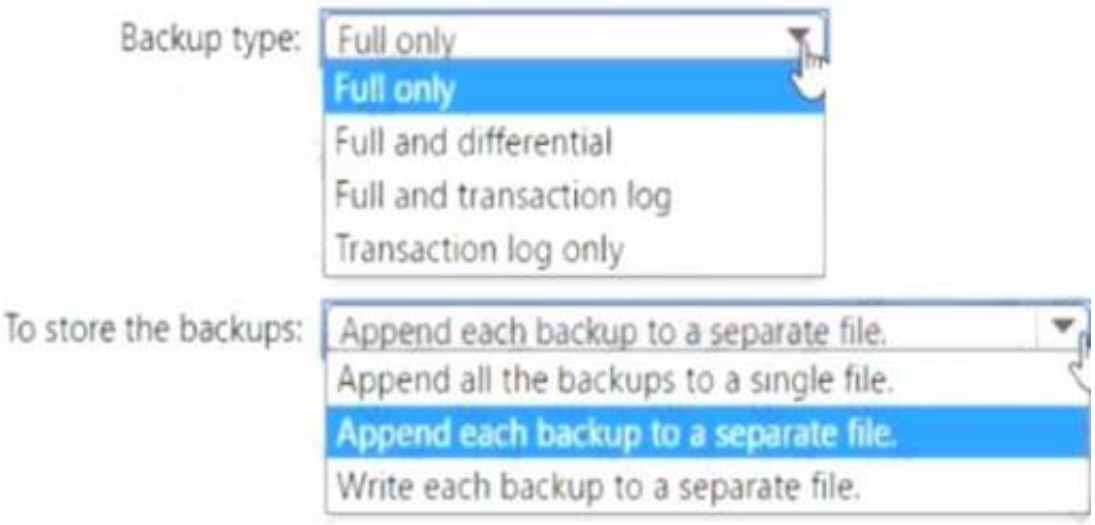

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

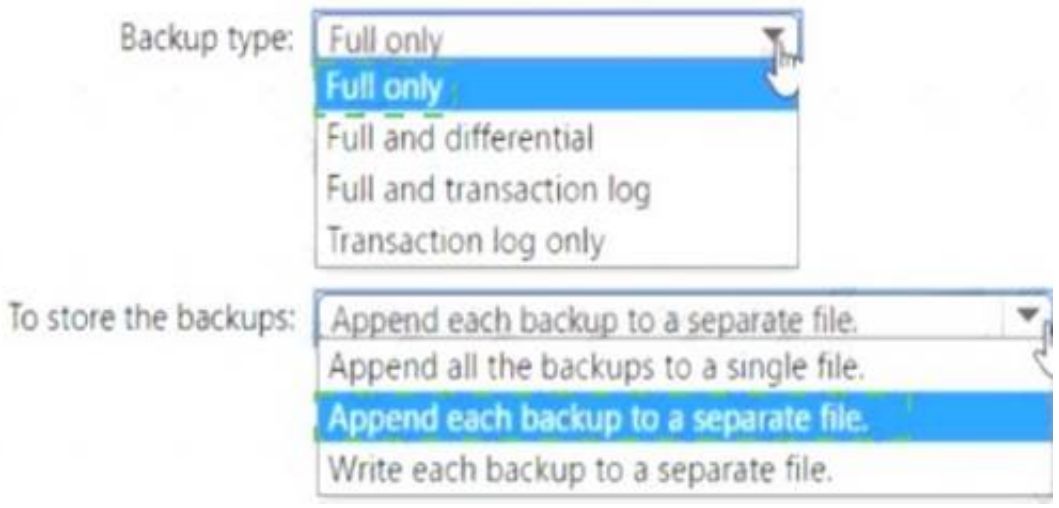

NEW QUESTION 216
- (Exam Topic 5)
You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.
You need to ensure that users from each company can view only the data of their respective company. Which two objects should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. a column encryption key
B. asymmetric keys
C. a function
D. a custom role-based access control (RBAC) role
E. a security policy

Answer: DE

Explanation:
Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes.
Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline

NEW QUESTION 220
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine.
You review the query plan shown in the following exhibit.

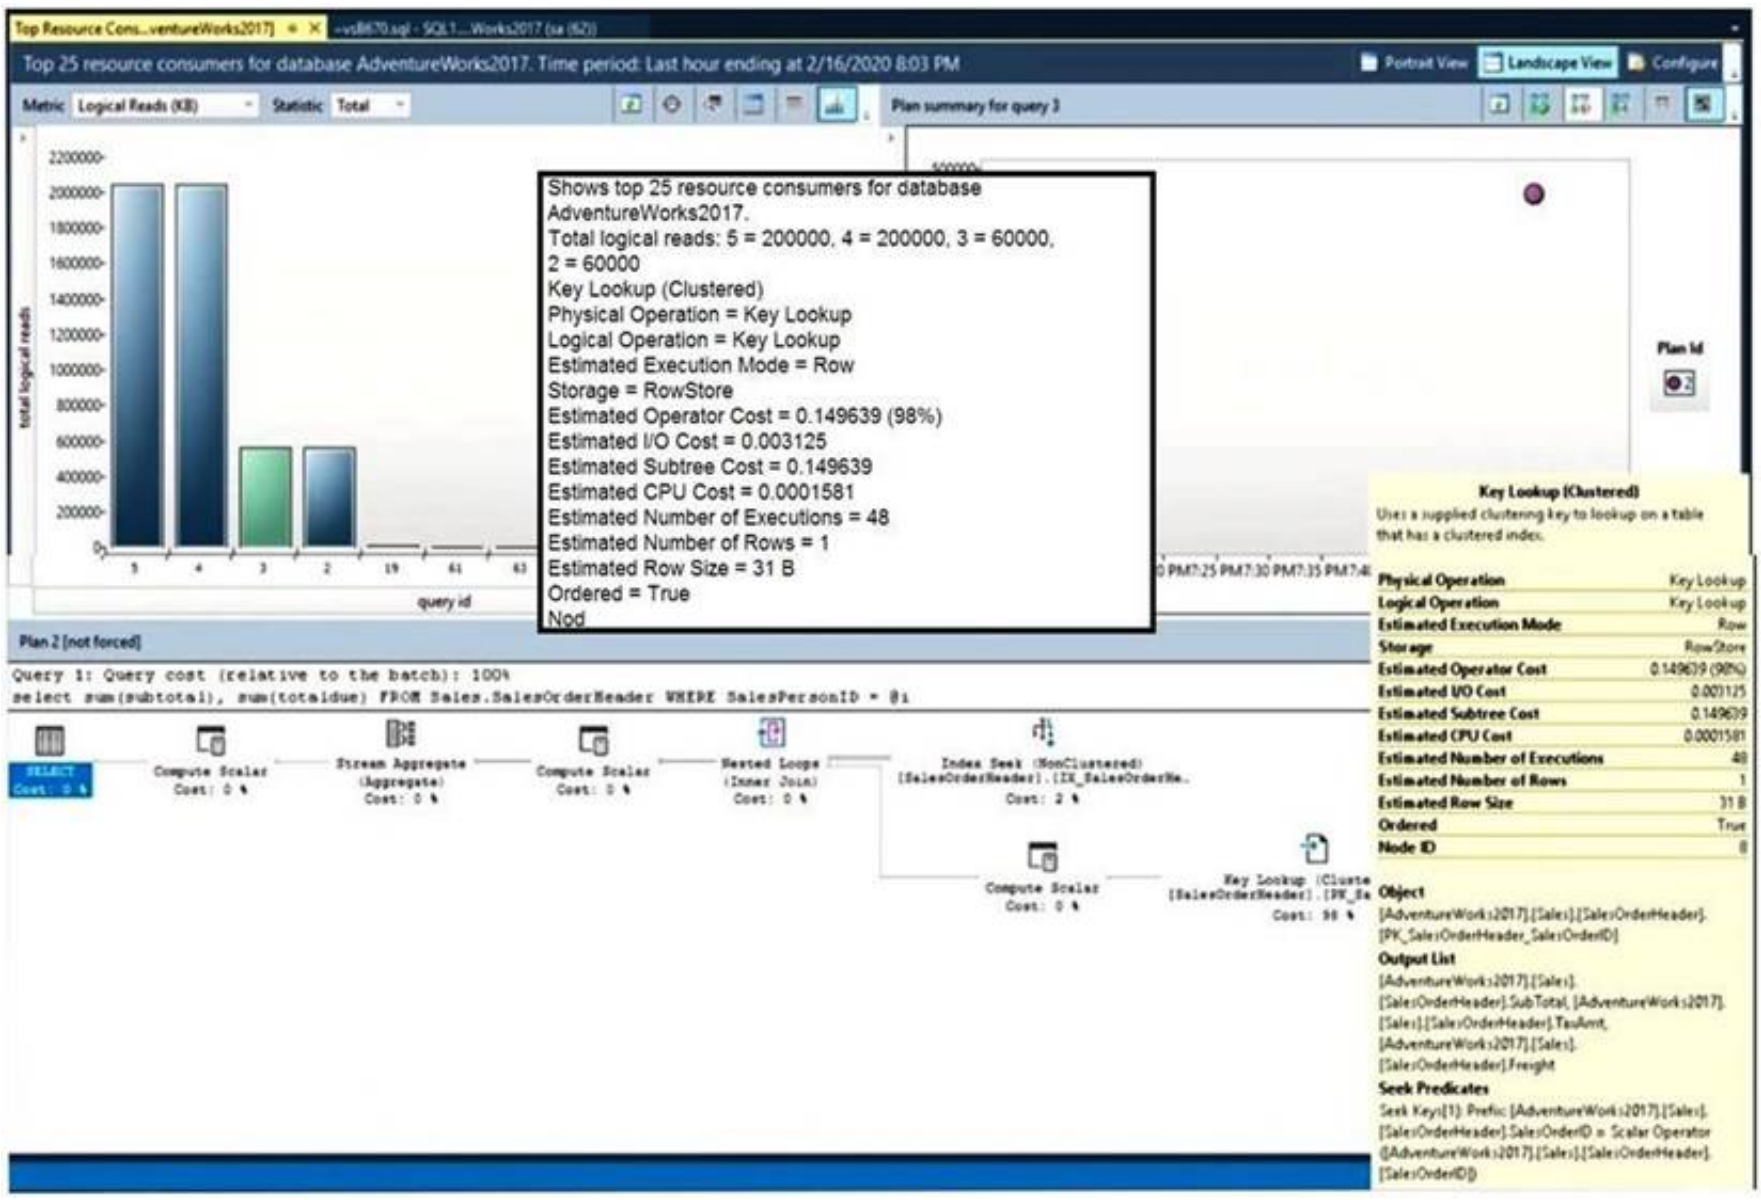

For each of the following statements, select yes if the statement is true. Otherwise, select no.
NOTE:Each correct selection is worth one point.

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	○	○
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	○	○
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	○	○

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 221
- (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 224
- (Exam Topic 5)
You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.
You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Encrypt the Salary column by using the randomized encryption type.
- Create a column encryption key.
- Enable Transparent Data Encryption (TDE).
- Encrypt the Salary column by using the deterministic encryption type.
- Apply a dynamic data mask to the Salary column.
- Create a column master key.

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create a column master key
Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.
Step 2: Create a column encryption key.
Step 3: Encrypt the Salary column by using the randomized encryption type.
Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.
Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine

NEW QUESTION 227
- (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 231
- (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.

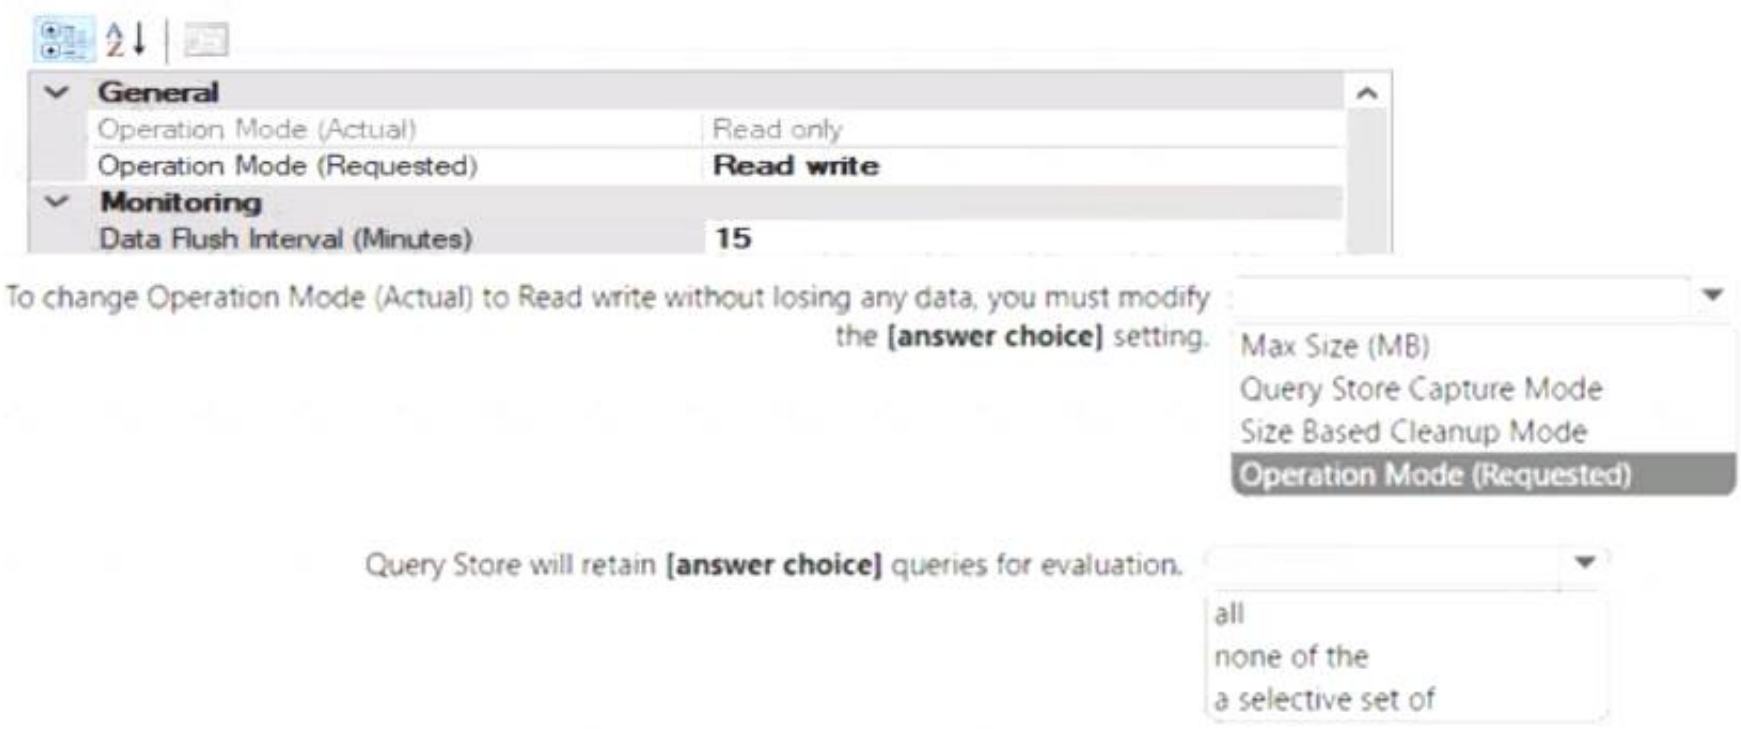

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text Description automatically generated

NEW QUESTION 234
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.
You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.
The on-premises network connects to VNET1 by using an ExpressRoute connection.
You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service. You need to configure VNET1 to support the migration.
What should you do?

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

Answer: A

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

NEW QUESTION 235
- (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 238
- (Exam Topic 5)
You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.
Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the
amount of time required to delete old data.
How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.
You may need to drag the split bar between panes or scroll to view content.
NOTE:Each correct selection is worth one point.

Values

CustomerKey
HASH
ROUND_ROBIN
REPLICATE
OrderDateKey
SalesOrderNumber

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: HASH
Box 2: OrderDateKey
In most cases, table partitions are created on a date column.
A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management. For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool

NEW QUESTION 242
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
SQL1	SQL Server on Azure Virtual Machines	Not applicable
db1	Microsoft SQL Server database	Hosted on SQL1
mysqlbackups	General purpose v2 storage account	Not applicable

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.
Which three Transact-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

```
RESTORE DATABASE db2 FROM URL = URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
     MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
     MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH RECOVERY,
     MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
     MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'

GO
```

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY
=
'sqlsamplebackup'
,SECRET = '<mystorageaccountaccesskey>';
GO
```

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-serv

NEW QUESTION 246
- (Exam Topic 5)
You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

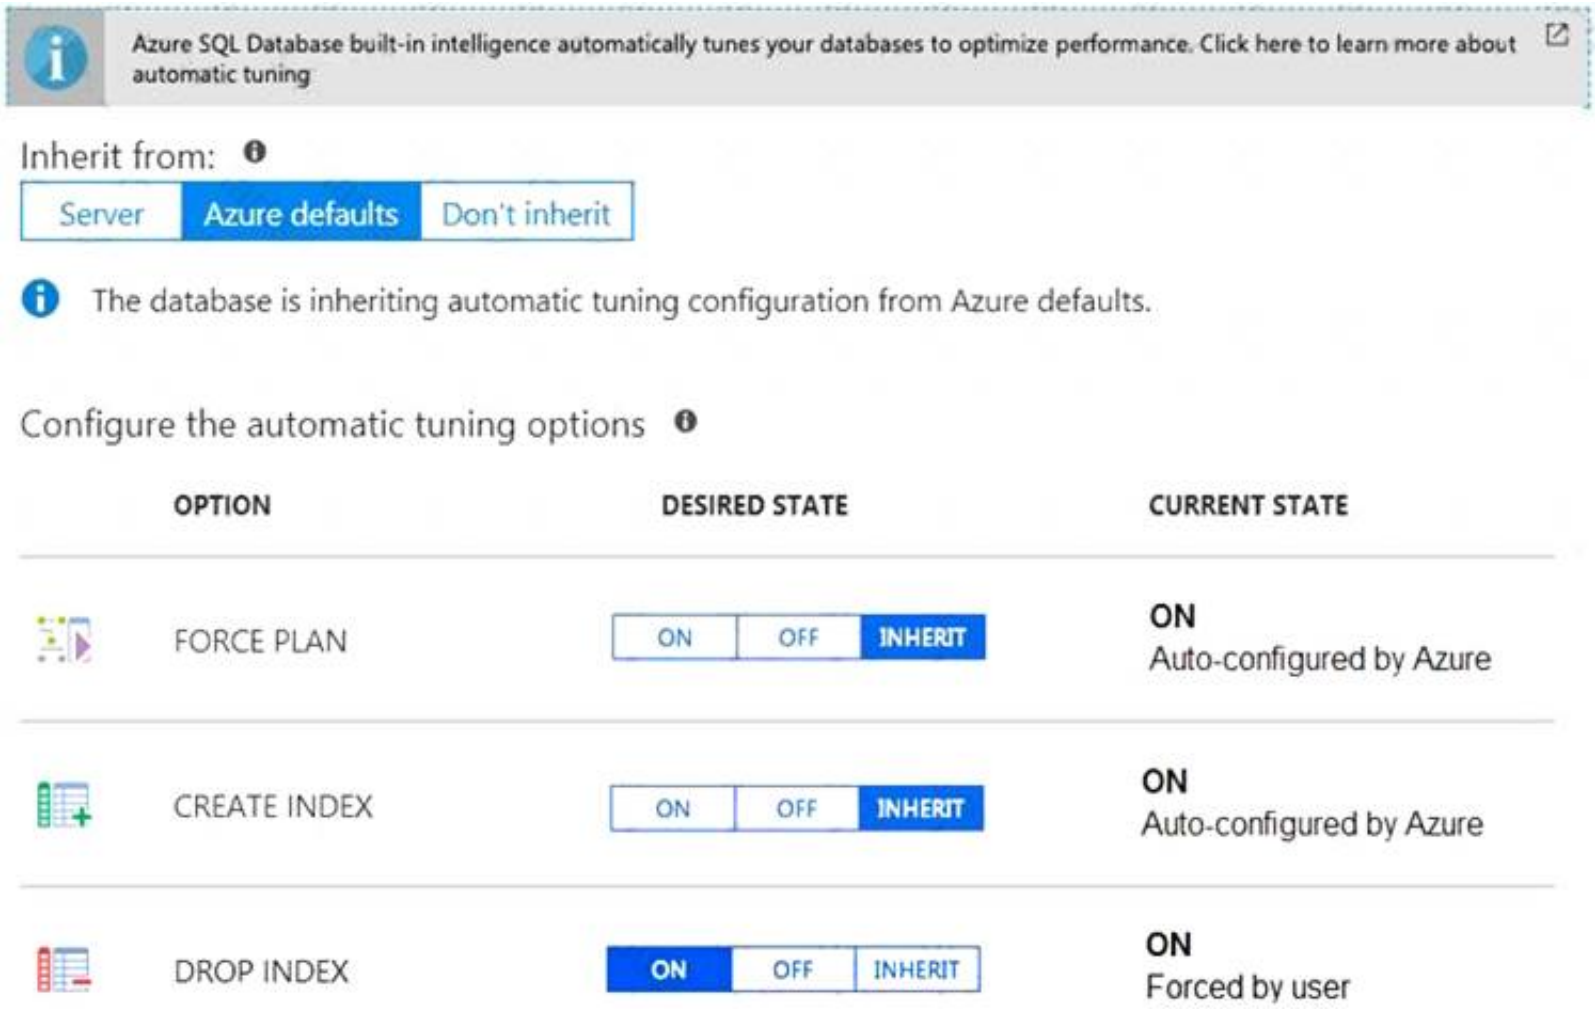

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE:Each correct selection is worth one point.

Statements	Yes	No
Nonclustered indexes will be added to tables to improve performance.	○	○
Columns will be added to existing indexes automatically.	○	○
The query execution plan will revert to a previous plan if query performance degrades.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
We see: Tuning option: Create index ON
CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.
Box 2: No
Box 3: Yes
FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

NEW QUESTION 249
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines named SQL1.
SQL1 contains an Extended Events session named session! that captures Microsoft SQL Server events. You need to correlate the session events with events captured by Event Tracing for Windows (ETW). What should you do for session1?

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

Answer: B

NEW QUESTION 251
- (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 254
- (Exam Topic 5)
You have an Azure virtual machine based on a custom image named VM1. VM1 hosts an instance of Microsoft SQL Server 2019 Standard.
You need to automate the maintenance of VM1 to meet the following requirements: Automate the patching of SQL Server and Windows Server.
Automate full database backups and transaction log backups of the databases on VM1. Minimize administrative effort.
What should you do first?

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

Answer: D

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

NEW QUESTION 258
- (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.

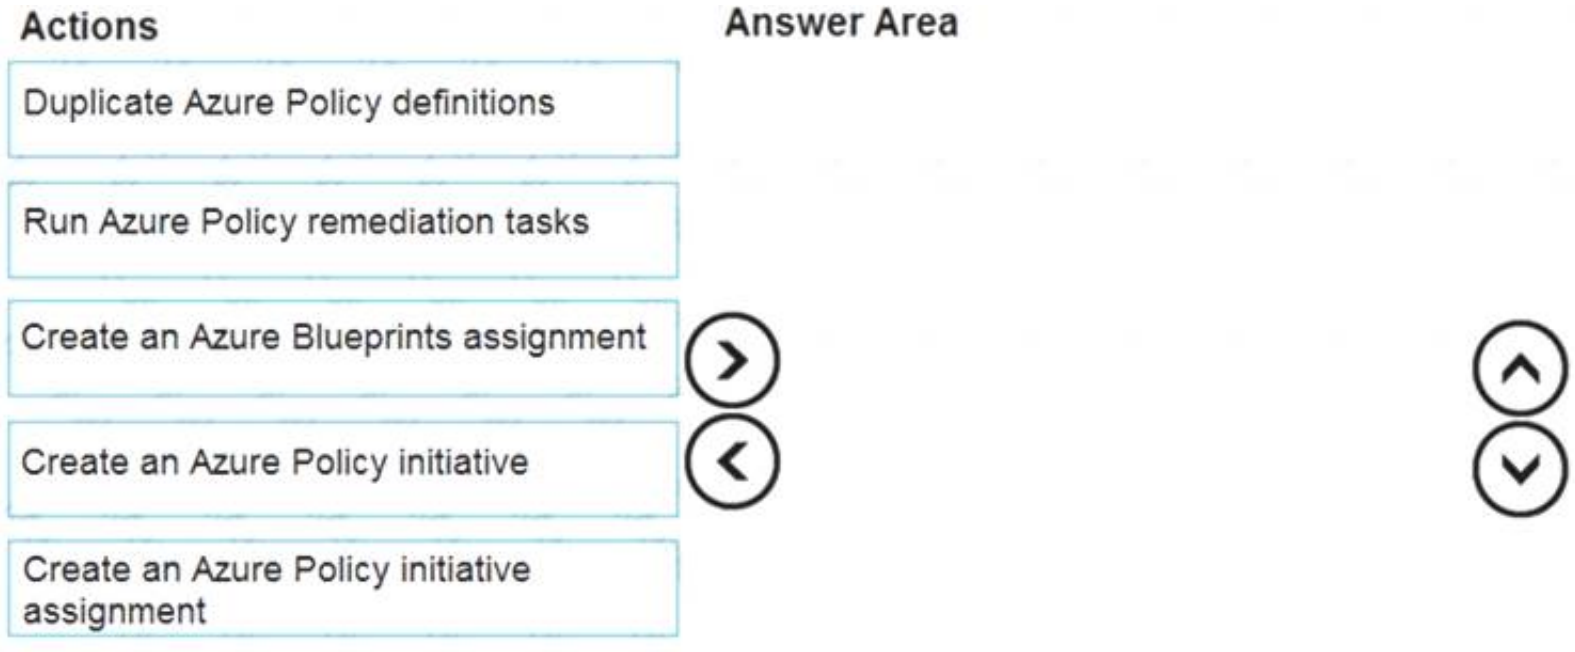

A. Mastered
B. Not Mastered

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 259
- (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 261
- (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.)

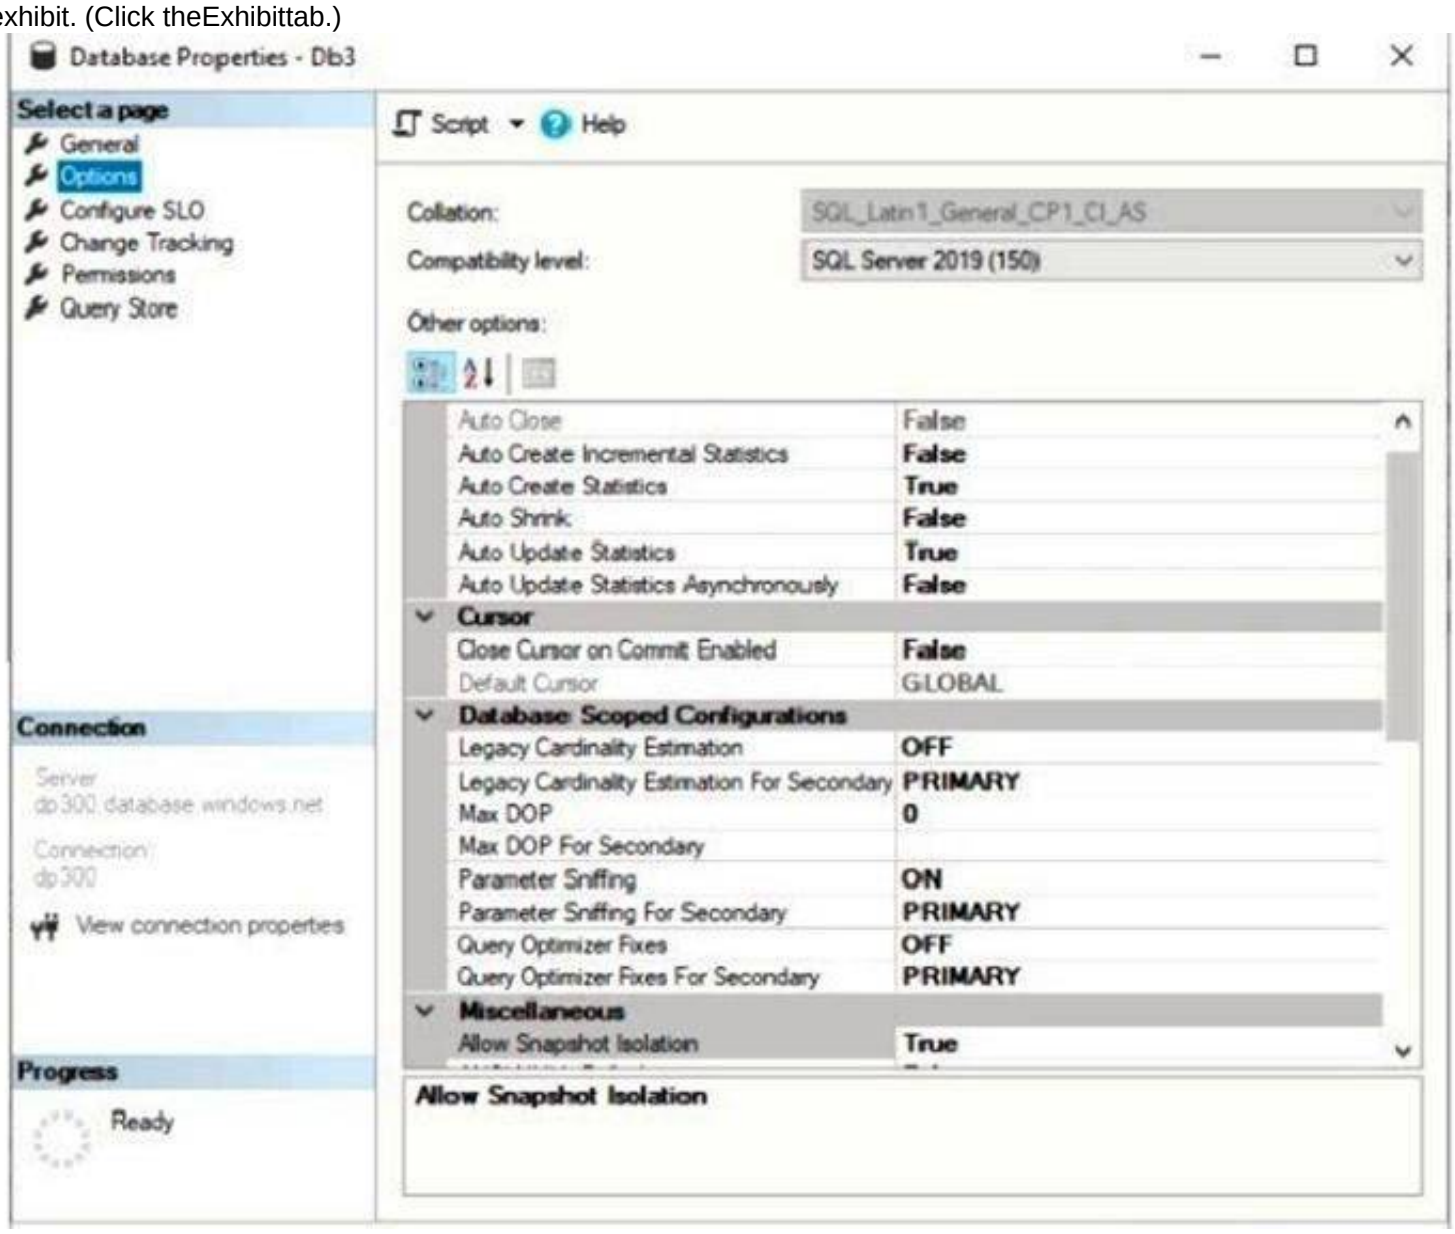

Which Transact-SQL command should you run?

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 262

- (Exam Topic 5)
You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.
You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication. You need to ensure that App1 can access DB2. The solution must meet the following requirements:

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

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

Answer: D

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

NEW QUESTION 265
- (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 267
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

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

You plan to migrate DB1 to an Azure SQL database.
What feature should be removed or replaced before DB1 can be migrated?

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

Answer: B

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

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

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

NEW QUESTION 268
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.
How should you complete the ARM template? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Answer Area

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, Word, email Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/create-template-quickstart?tabs=azure-powe

NEW QUESTION 271
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to encrypt DB1. The solution must meet the following requirements;
• Encrypt data in motion.
• Support comparison operators.
• Provide randomized encryption.
What should you include in the solution?

A. Always Encrypted
B. column-level encryption
C. Transparent Data Encryption (TDE)
D. Always Encrypted with secure enclaves

Answer: A

NEW QUESTION 276
- (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 280
- (Exam Topic 5)
You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables.
Which distribution type should you recommend to minimize data movement?

A. HASH
B. REPLICATE
C. ROUND_ROBIN

Answer: B

Explanation:
A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication requires extra storage, though, and isn't practical for large tables.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu

NEW QUESTION 284
- (Exam Topic 5)
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

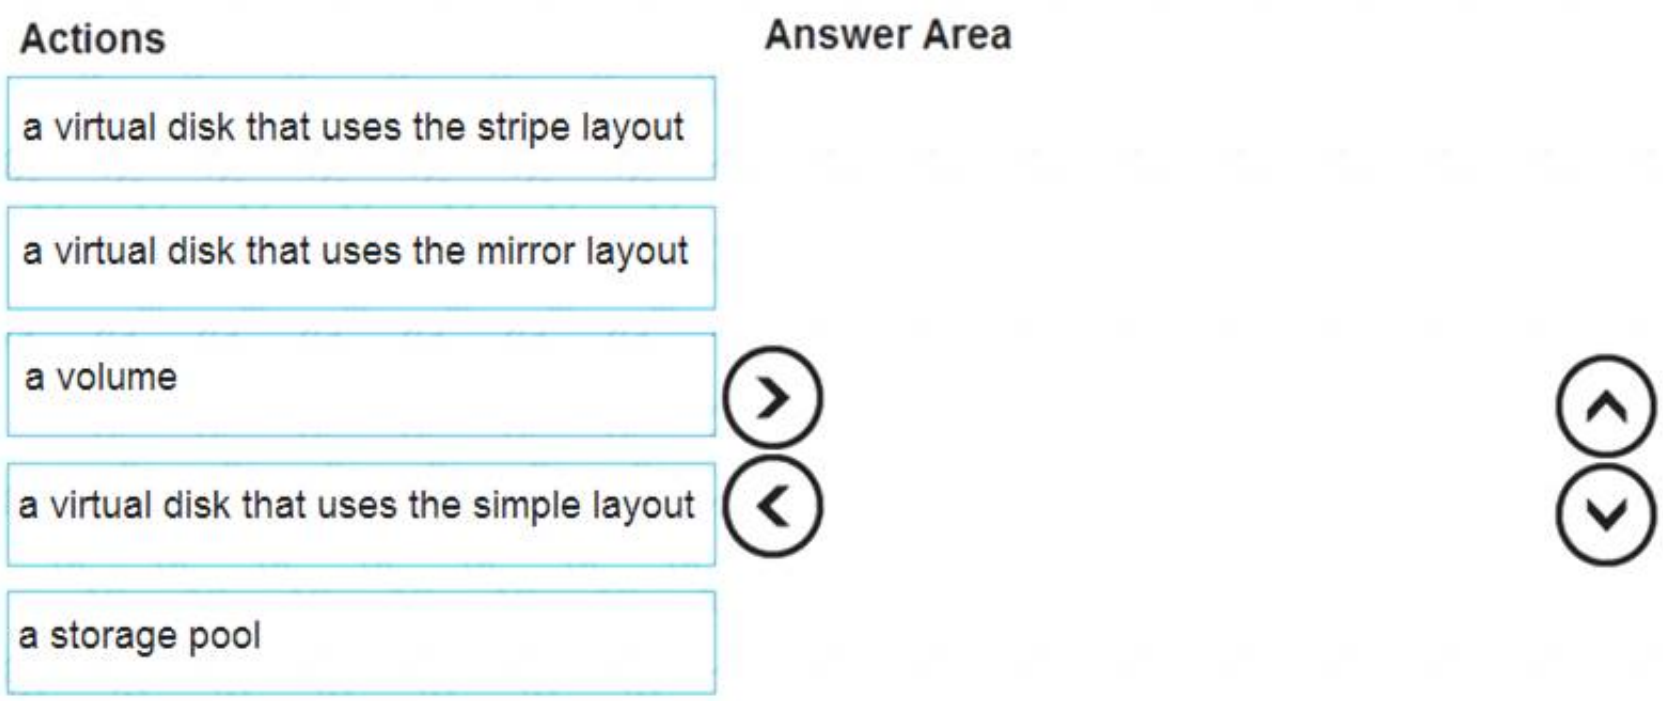

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Follow these same steps to create striped virtual disk:
- Create Log Storage Pool.
- Create Virtual Disk
- Create Volume

Box 1: a storage pool
Box 2: a virtual disk that uses stripe layout
Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.
Box 3: a volume Reference:
https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/

NEW QUESTION 288
- (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 290
- (Exam Topic 5)
You have an Azure Data Factory pipeline that is triggered hourly. The pipeline has had 100% success for the past seven days.
The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

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

What is a possible cause of the error?

A. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.
B. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.
C. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.
D. The pipeline was triggered too early.

Answer: B

NEW QUESTION 291

- (Exam Topic 5)
You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:
- Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: B

Explanation:
Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.
Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.
Reference:
https://docs.microsoft.com/en-us/azure/private-link/private-link-overview

NEW QUESTION 296
- (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 299
- (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.

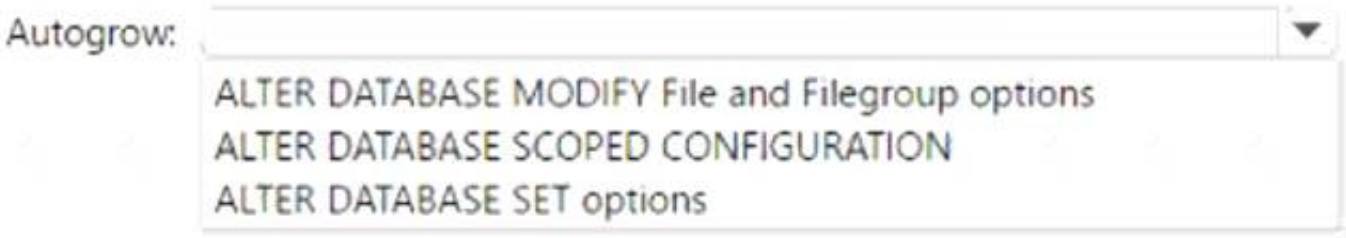

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 304
- (Exam Topic 5)
You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.
You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.
Which resource type should you use to create the alert?

A. Resource Groups
B. SQL Servers
C. SQL Databases

D. SQL Virtual Machines

Answer: C

Explanation:
There are resource types related to application code, compute infrastructure, networking, storage + databases. You can deploy up to 800 instances of a resource type in each resource group.
Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types

NEW QUESTION 306
- (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 310
- (Exam Topic 5)
Your company uses Azure Stream Analytics to monitor devices.
The company plans to double the number of devices that are monitored.
You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.
Which metric should you monitor?

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

Answer: D

Explanation:
The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.
The watermark delay metric can rise due to:
* 1. Not enough processing resources in Stream Analytics to handle the volume of input events.
* 2. Not enough throughput within the input event brokers, so they are throttled.
* 3. Output sinks are not provisioned with enough capacity, so they are throttled. Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling

NEW QUESTION 311
- (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 316
- (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.

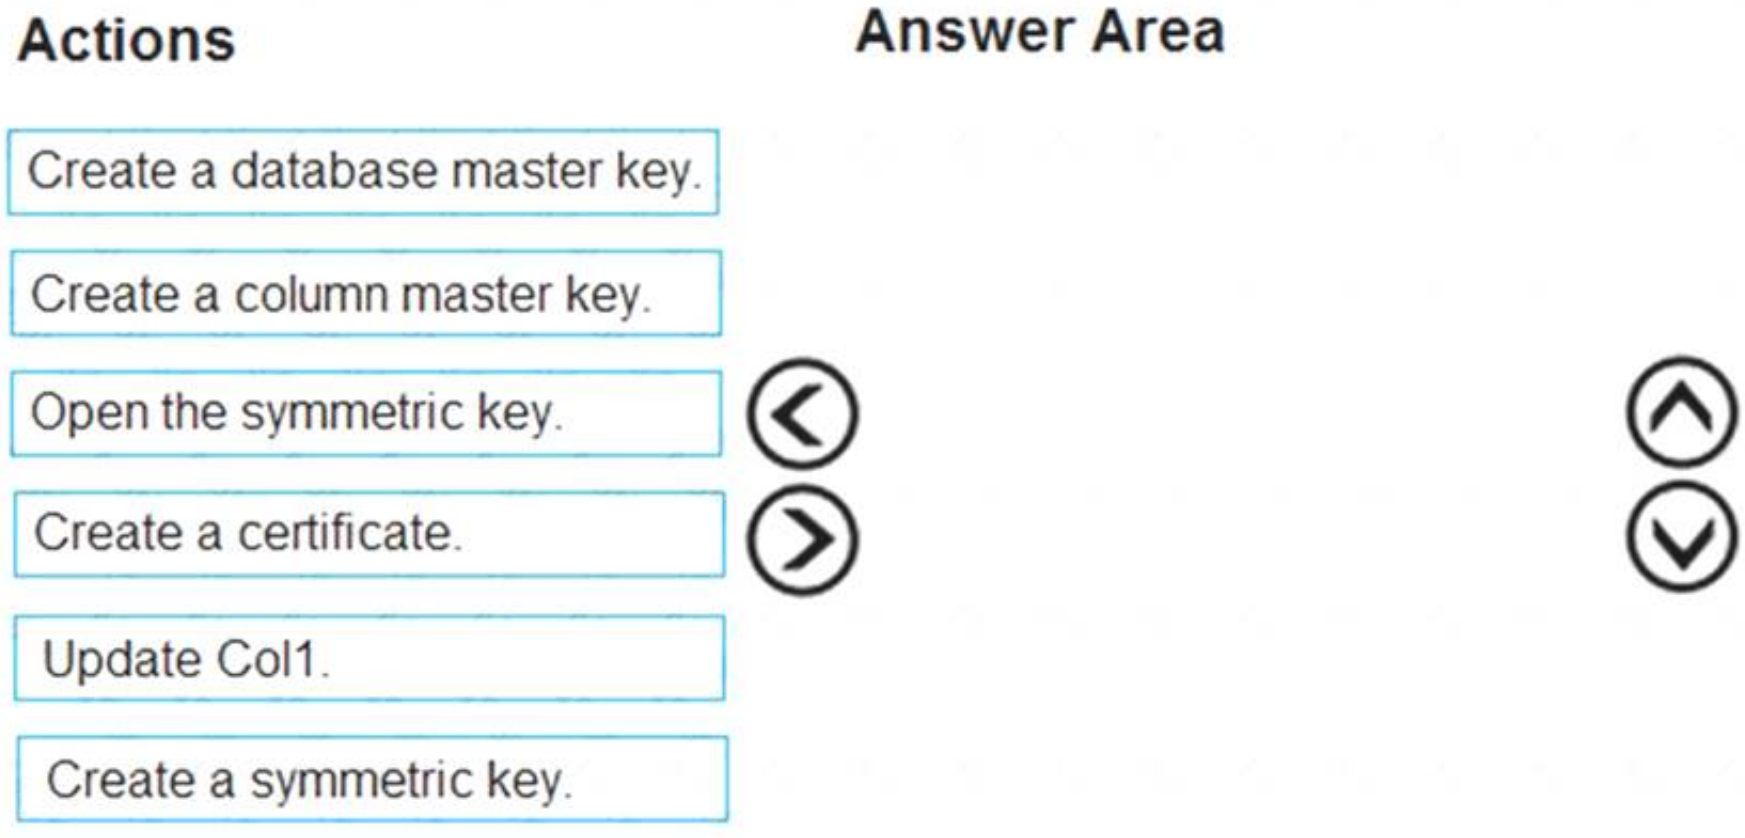

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 317
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1 and an Azure Storage account named storage1.
You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:
* Use block blob storage.
* Maximize security.

A. Generate a shared access signature (SAS)
B. Enable infrastructure encryption.
C. Create an access policy.
D. Rotate the storage keys

Answer: B

NEW QUESTION 321
- (Exam Topic 5)
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?

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

Answer: A

Explanation:
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

NEW QUESTION 324
- (Exam Topic 5)
You have a Microsoft SQL Server database named DB1 that contains a table named Table1. The database role membership for a user named User1 is shown in the following exhibit.

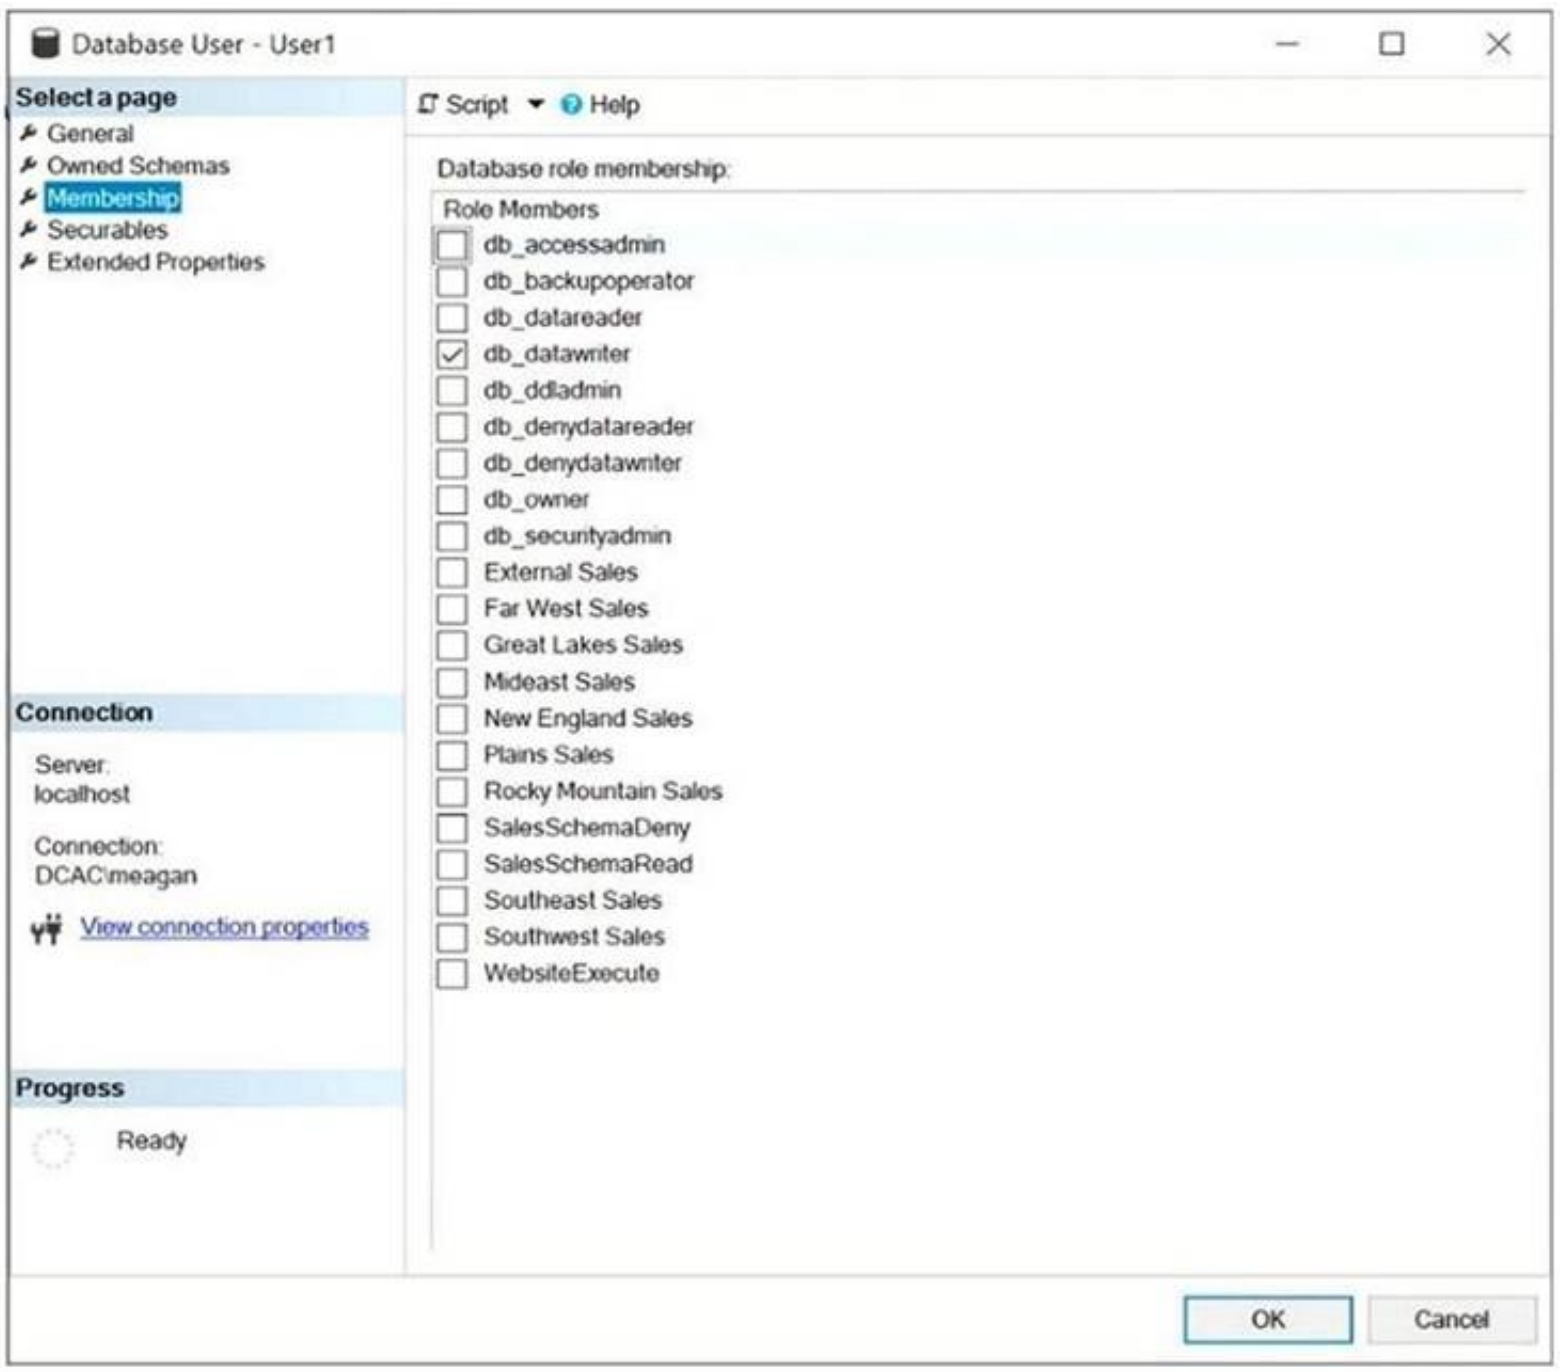

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE:Each correct selection is worth one point.

User1 can **[answer choice]**.

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

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

- db_datareader
- db_ddladmin
- db_denydatareader
- db_denydatawriter

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: delete a row from Table1
Members of the db_datawriter fixed database role can add, delete, or change data in all user tables. Box 2: db_datareader
Members of the db_datareader fixed database role can read all data from all user tables. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles

NEW QUESTION 325
- (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 329
- (Exam Topic 5)
You have a 50-TB Microsoft SQL Server database named DB1.
You need to reduce the time it takes to perform database consistency checks of DB1.
Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

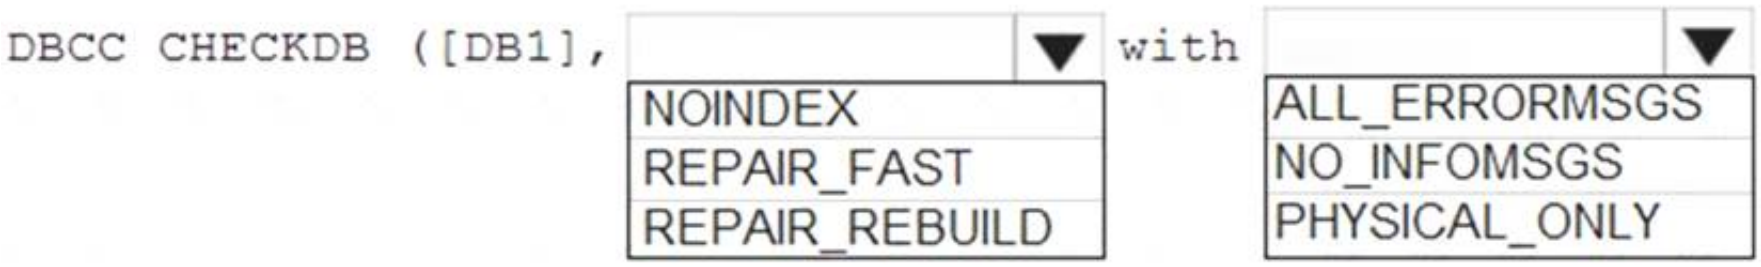

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Table Description automatically generated with low confidence
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser

NEW QUESTION 334
- (Exam Topic 5)
You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.
During peak usage, the database will require the following:
- 24 cores
- 500 GB of storage
- 124 GB of memory
- More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard. Which service tier should you use during peak usage?

A. Business Critical
B. Premium
C. Hyperscale

Answer: A

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

NEW QUESTION 338
- (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 339
- (Exam Topic 5)
You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:
- Support failovers that do not require client applications to change their connection strings.

- Replicate the database to a secondary Azure region.
- Support failover to the secondary region. What should you include in the recommendation?

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

Answer: A

Explanation:
Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview

NEW QUESTION 343
- (Exam Topic 4)
You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.
Which Azure Storage functionality should you include in the solution?

A. time-based retention
B. change feed
C. lifecycle management
D. soft delete

Answer: C

Explanation:
The lifecycle management policy lets you:
Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts

NEW QUESTION 345
- (Exam Topic 3)
Which windowing function should you use to perform the streaming aggregation of the sales data?

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

Answer: D

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

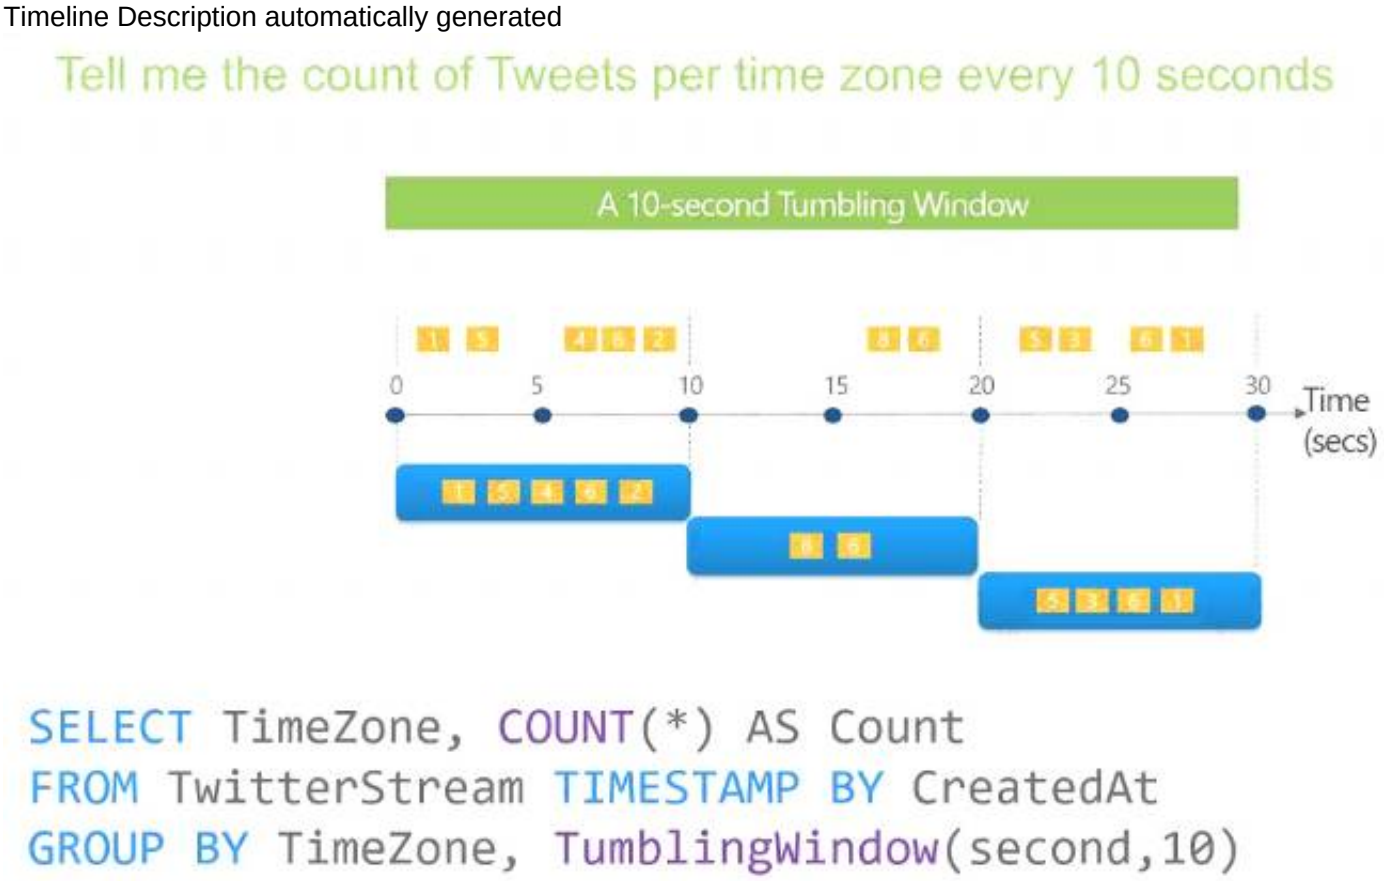

Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 350

- (Exam Topic 2)
Which audit log destination should you use to meet the monitoring requirements?

A. Azure Storage
B. Azure Event Hubs
C. Azure Log Analytics

Answer: C

Explanation:
Scenario: Use a single dashboard to review security and audit data for all the PaaS databases.
With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.
Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards

NEW QUESTION 351
- (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.

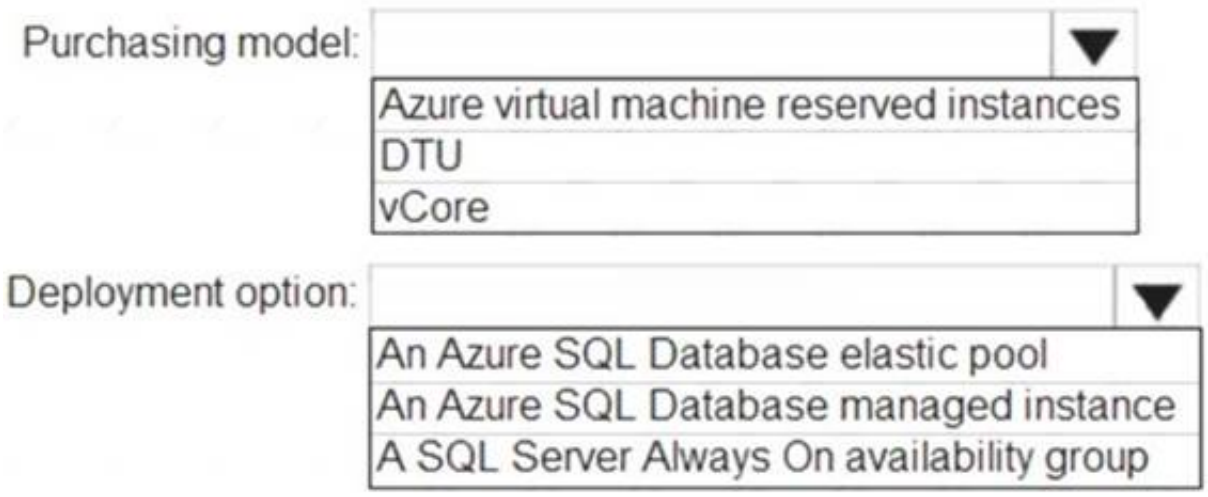

A. Mastered
B. Not Mastered

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 354
- (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.

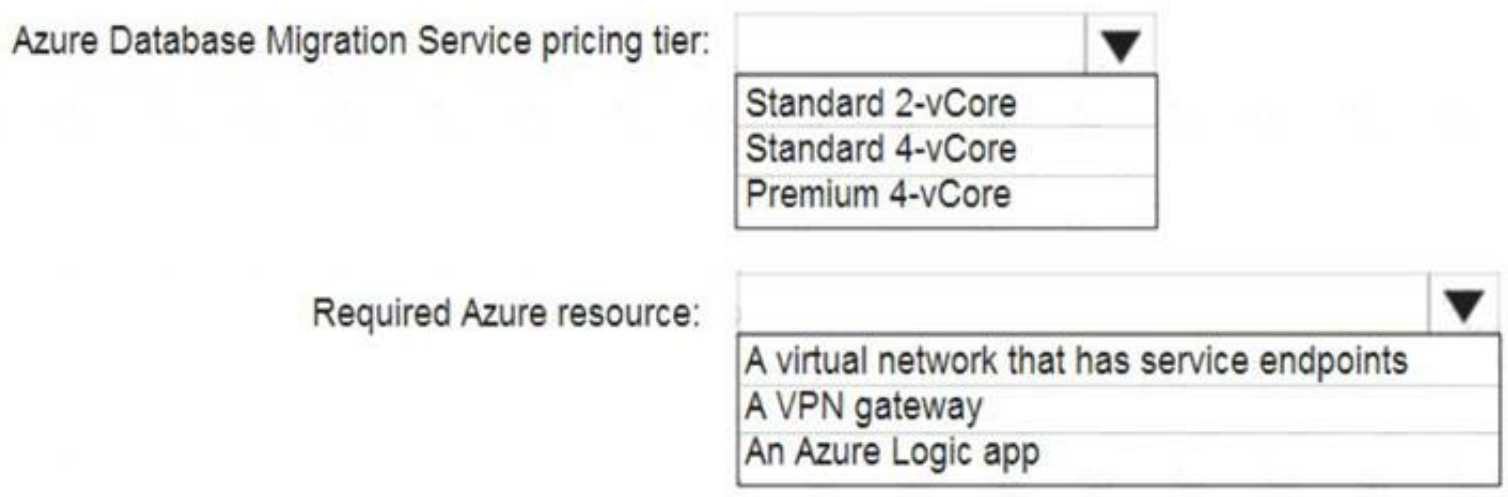

A. Mastered
B. Not Mastered

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 358
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-300 Practice Exam Features:

* DP-300 Questions and Answers Updated Frequently

* DP-300 Practice Questions Verified by Expert Senior Certified Staff

* DP-300 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-300 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The DP-300 Practice Test Here