Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

NEW QUESTION 1
- (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.

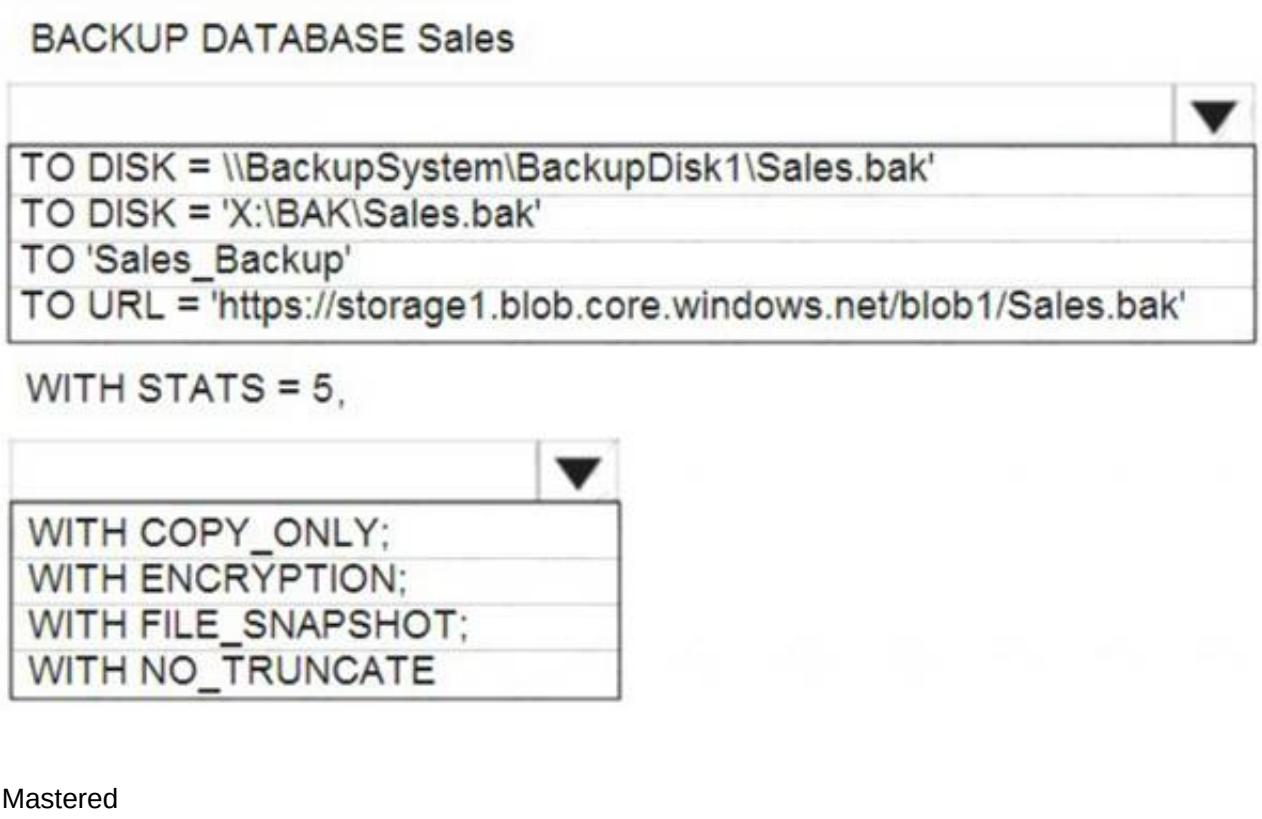

A. Mastered
B. Not Mastered

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 2
- (Exam Topic 5)
You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases. The databases perform slower than expected.
You need to identify whether the performance issue relates to the use of tempdb on sqlsrv1. What should you do?

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

Answer: D

Explanation:
The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting.
You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-use-diagnostics-log

NEW QUESTION 3
- (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 the select * from sys.database_scoped_configurations Transact-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 4
- (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.

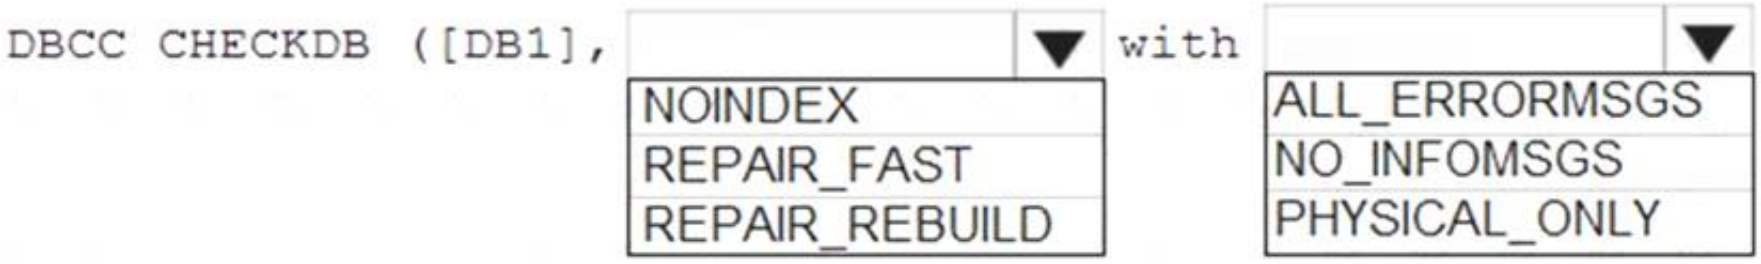

A. Mastered
B. Not Mastered

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

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

Answer: C

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

NEW QUESTION 6
- (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 7
- (Exam Topic 5)
You have an Azure SQL database named DB1. You run a query while connected to DB1.

You review the actual execution plan for the query, and you add an index to a table referenced by the query. You need to compare the previous actual execution plan for the query to the Live Query Statistics.
What should you do first in Microsoft SQL Server Management Studio (SSMS)?

A. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
B. Run the SET SHOWPLAN_ALL Transact-SQL statement.
C. Save the actual execution plan.
D. Enable Query Store for DB1.

Answer: C

Explanation:
The Plan Comparison menu option allows side-by-side comparison of two different execution plans, for easier identification of similarities and changes that explain the different behaviors for all the reasons stated above. This option can compare between:
Two previously saved execution plan files (.sqlplan extension).
One active execution plan and one previously saved query execution plan. Two selected query plans in Query Store.

NEW QUESTION 8
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.
You need to resolve the performance issues. What should you do?

A. Implement memory-optimized tables.
B. Run the dbcc flushprocindB command.
C. Replace the sequential index keys with nonsequential keys.
D. Run the dbcc dbreindex command.

Answer: A

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

NEW QUESTION 9
- (Exam Topic 5)
You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in fiveminute intervals and report only events that arrive during the interval.
The output will be sent to a Delta Lake table. Which output mode should you use?

A. complete
B. append
C. update

Answer: A

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

NEW QUESTION 10
- (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: B

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 10
- (Exam Topic 5)
You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.
Several users execute adhoc queries to DW1 concurrently. You regularly perform automated data loads to DW1.
You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.
What should you do?

A. Assign a smaller resource class to the automated data load queries.

B. Create sampled statistics to every column in each table of DW1.
C. Assign a larger resource class to the automated data load queries.
D. Hash distribute the large fact tables in DW1 before performing the automated data loads.

Answer: C

Explanation:
The performance capacity of a query is determined by the user's resource class.
Smaller resource classes reduce the maximum memory per query, but increase concurrency. Larger resource classes increase the maximum memory per query, but reduce concurrency. Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadman

NEW QUESTION 15
- (Exam Topic 5)
You deploy a database to an Azure SQL Database managed instance.
You need to prevent read queries from blocking queries that are trying to write to the database. Which database option should set?

A. PARAMETERIZATION to FORCED
B. PARAMETERIZATION to SIMPLE
C. Delayed Durability to Forced
D. READ_COMMITTED_SNAPSHOT to ON

Answer: D

Explanation:
In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:
- The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.
- The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql

NEW QUESTION 17
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.
You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders.
What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

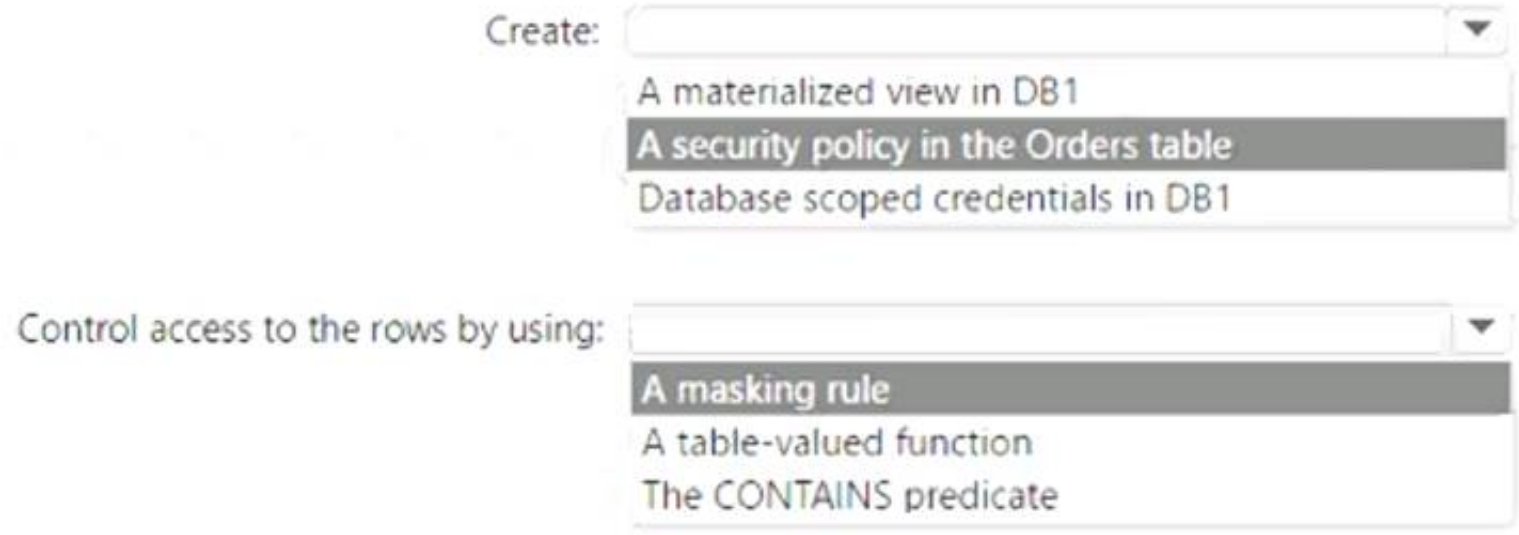

A. Mastered
B. Not Mastered

Answer: A

Explanation:

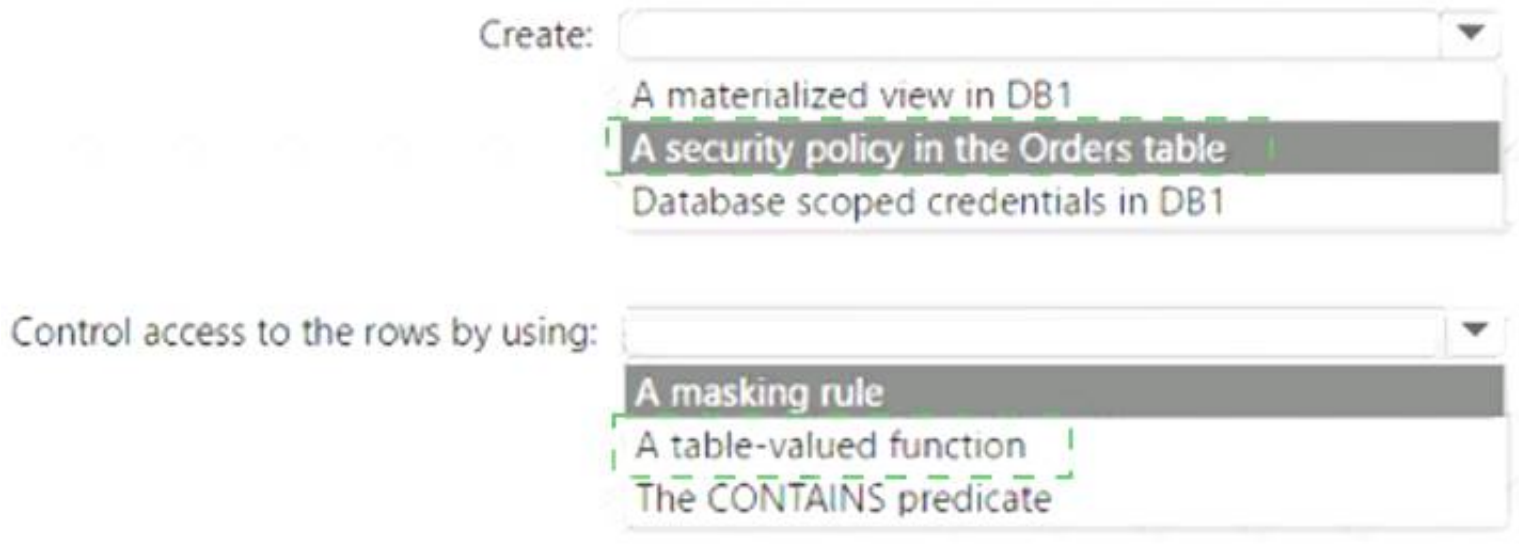

NEW QUESTION 21
- (Exam Topic 5)
You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.
You need to ensure that an automatic email notification is sent once the job completes. What should you include in the solution?

A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
B. From SQL Server Management Studio (SSMS), run sp_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 24
- (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 25
- (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 27
- (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.

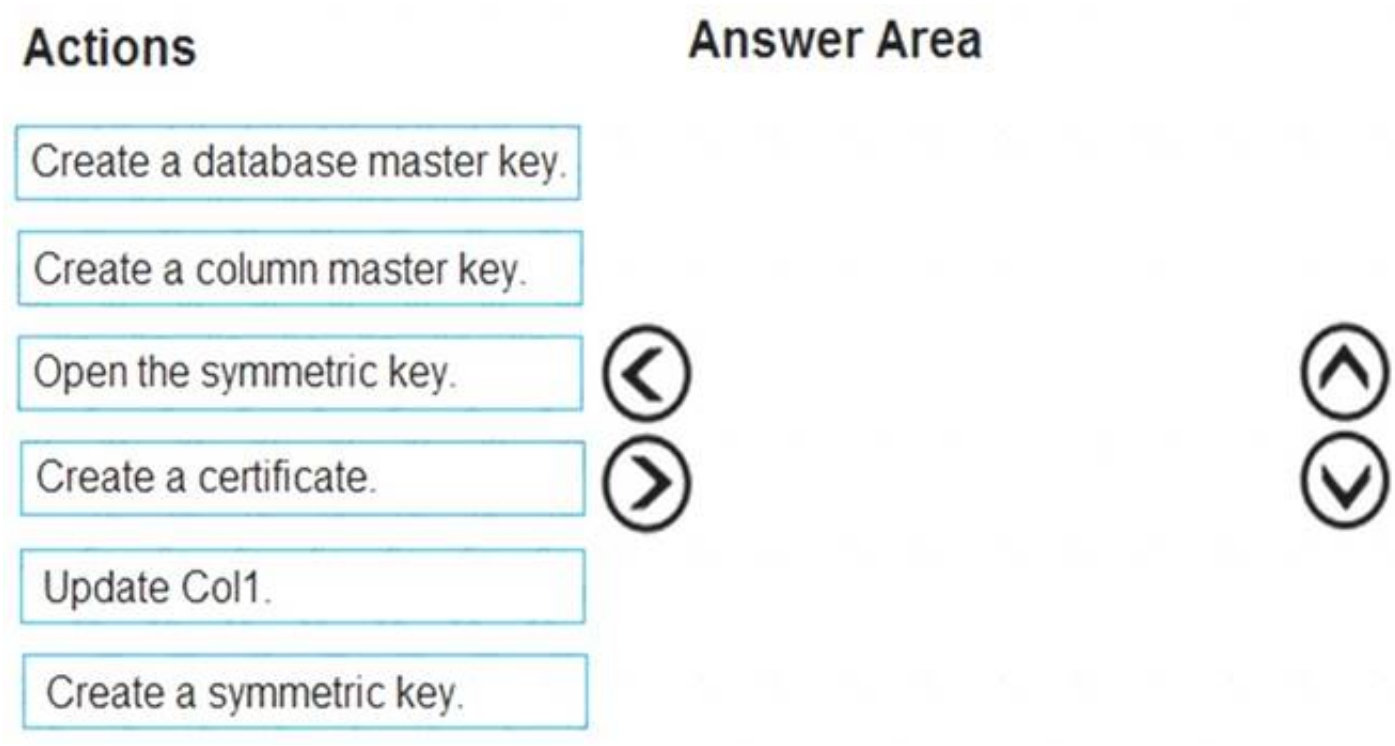

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 29
- (Exam Topic 5)
You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the Exhibit tab.)

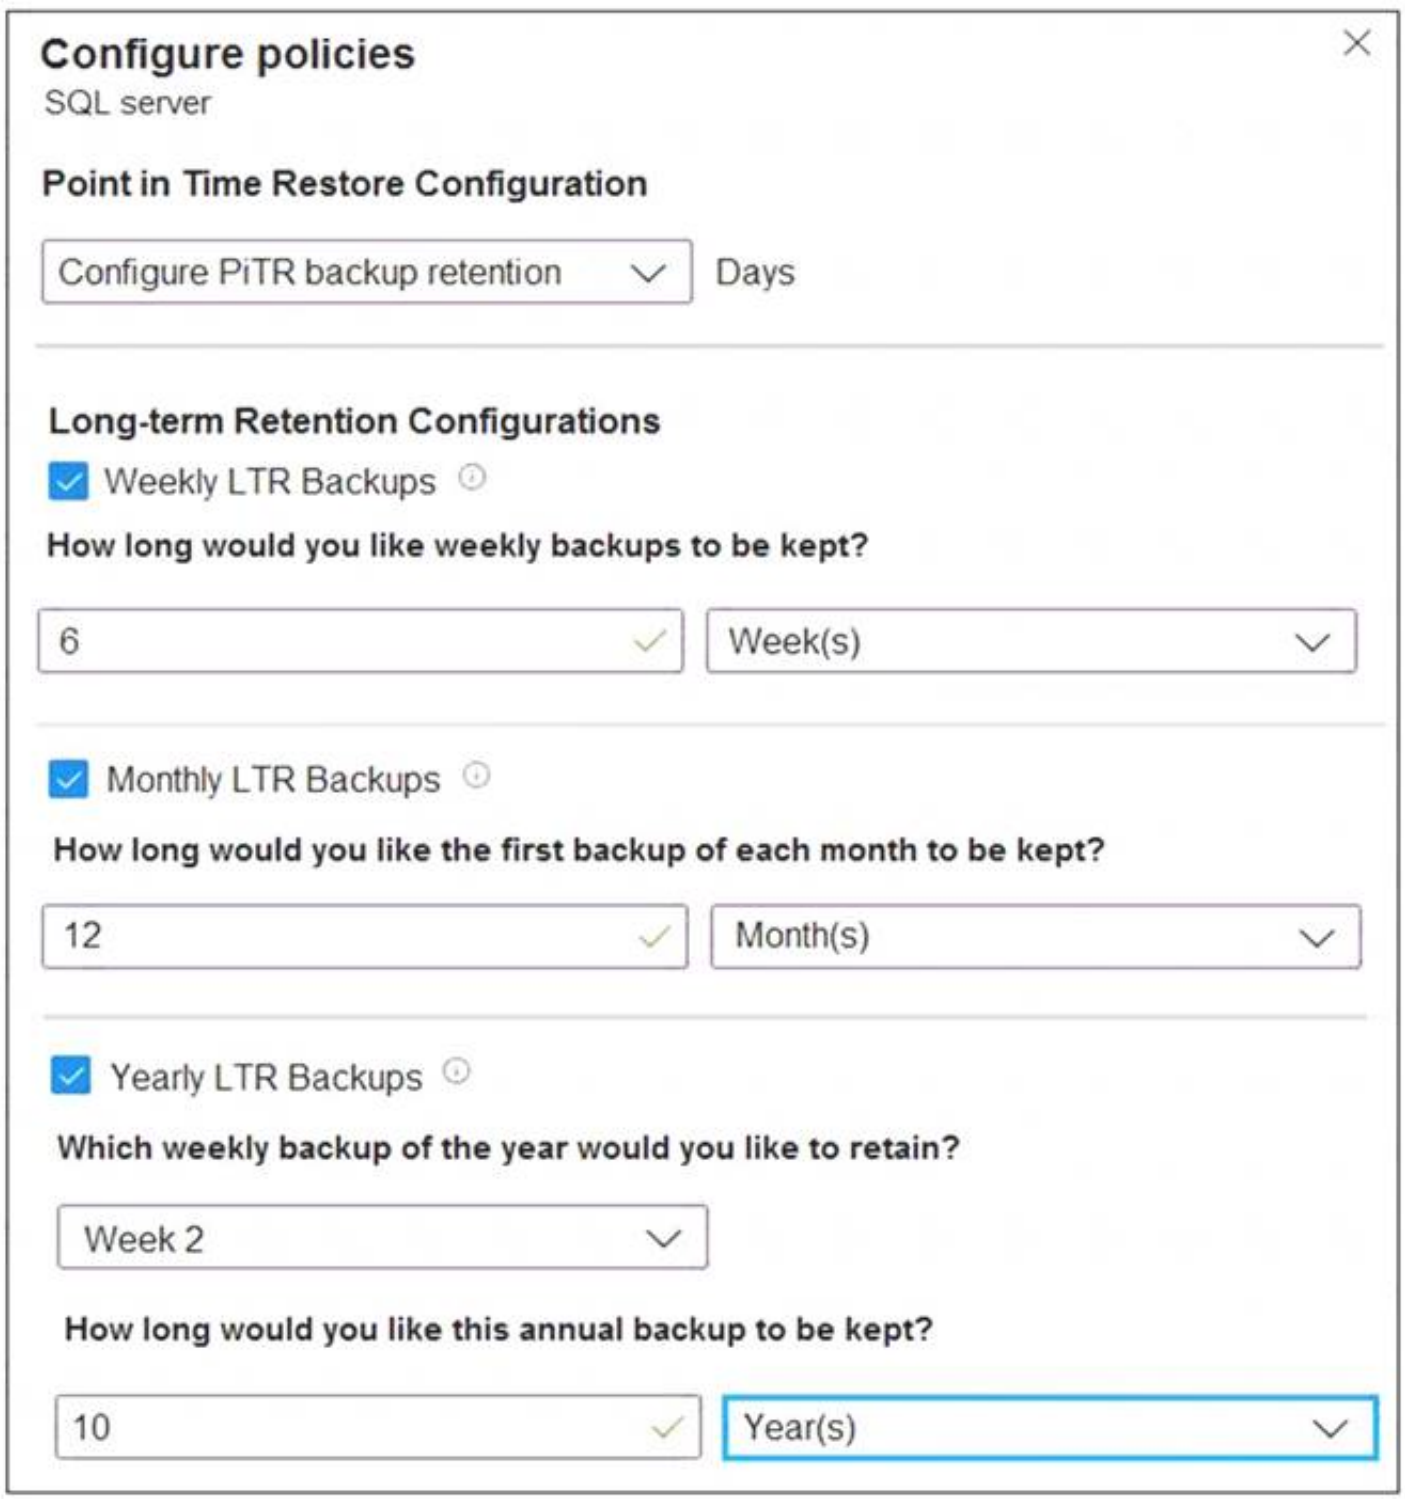

The first weekly backup occurred on January 4, 2020. The dates for the first 10 weekly backups are:

- January 4, 2020
- January 11, 2020
- January 18, 2020
- January 25, 2020
- February 1, 2020
- February 8, 2020
- February 15, 2020
- February 22, 2020
- February 29, 2020
- March 7, 2020

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.

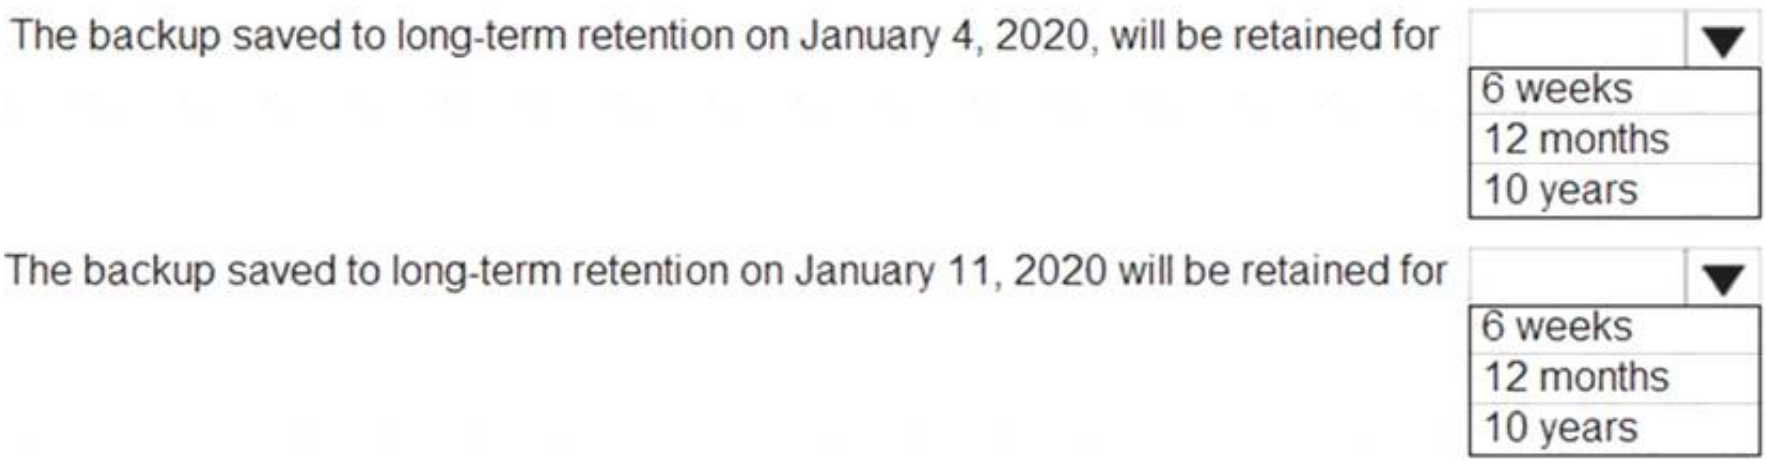

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 33
- (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 36
- (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.

Answer Area

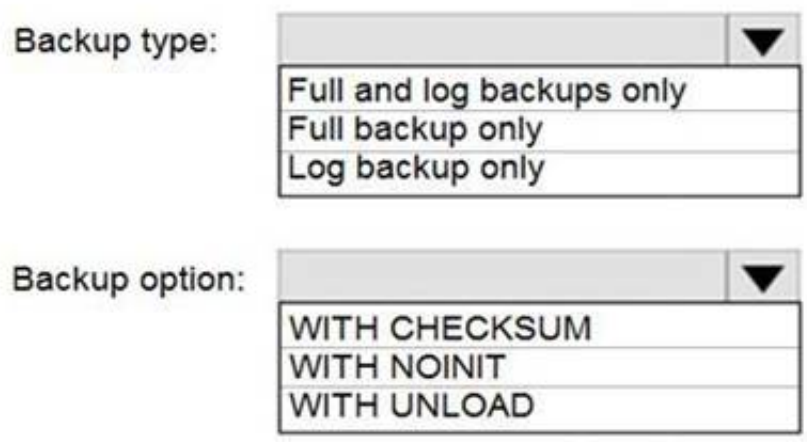

A. Mastered
B. Not Mastered

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 41
- (Exam Topic 5)
You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

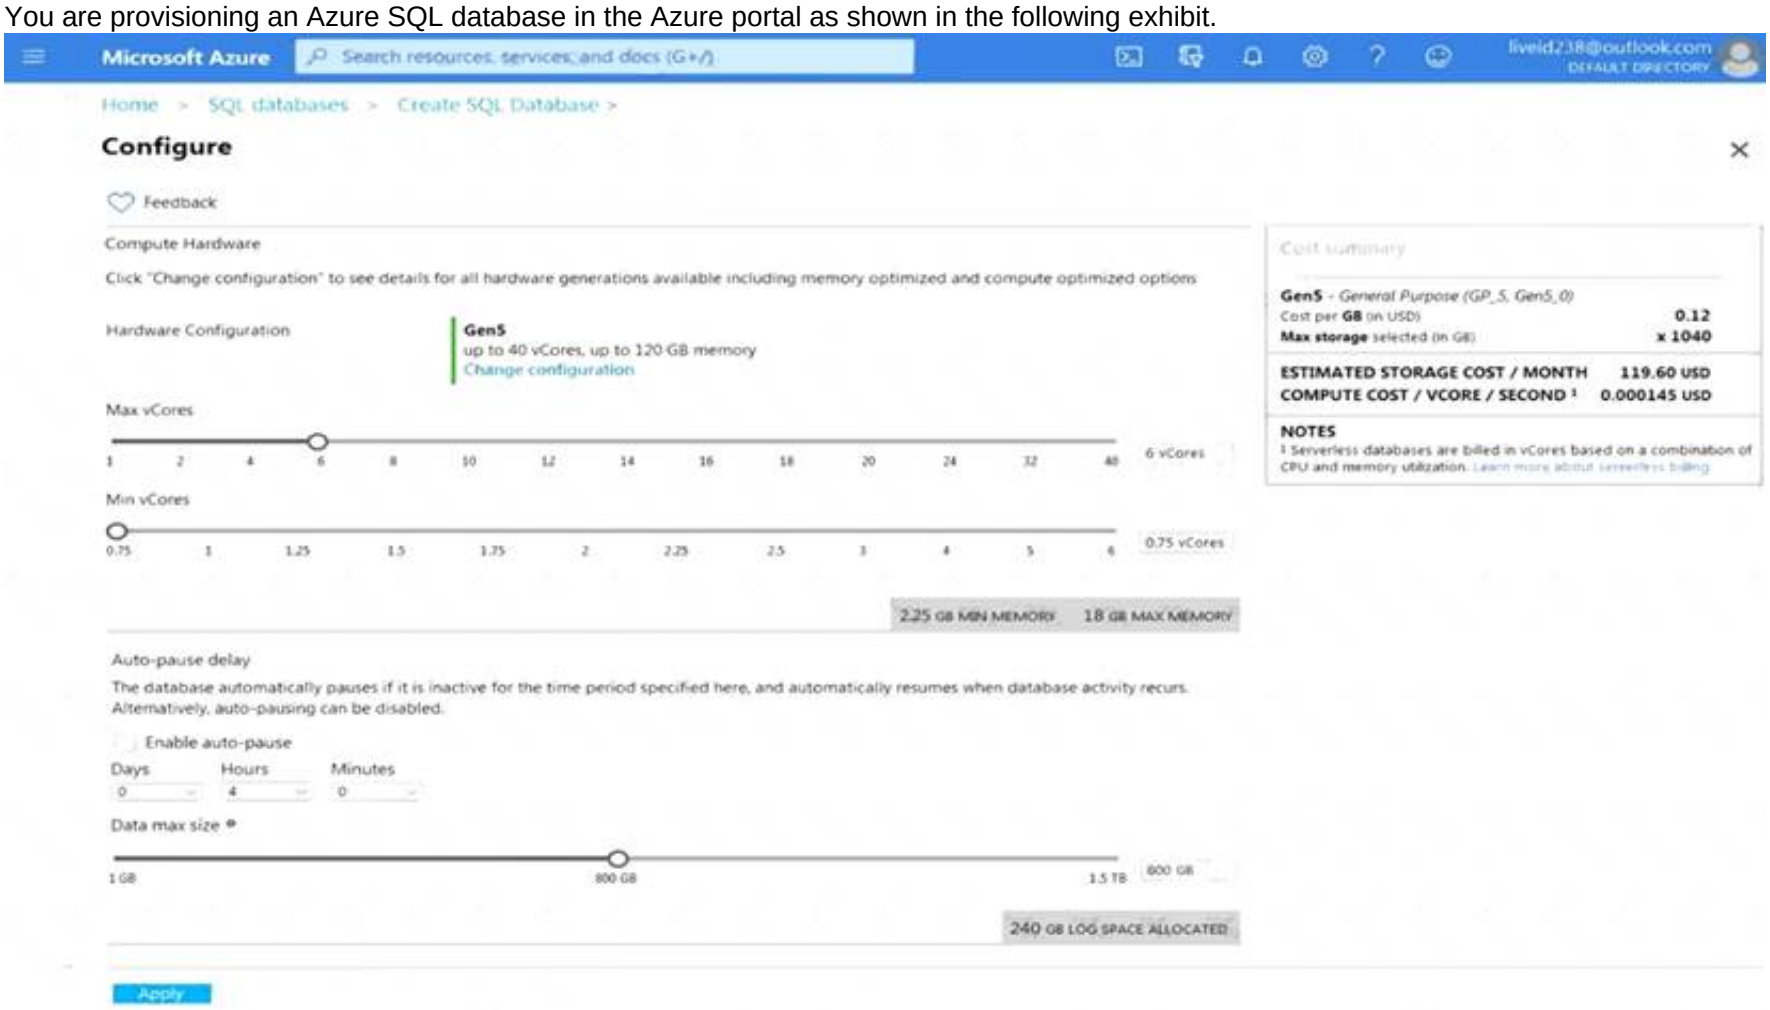

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

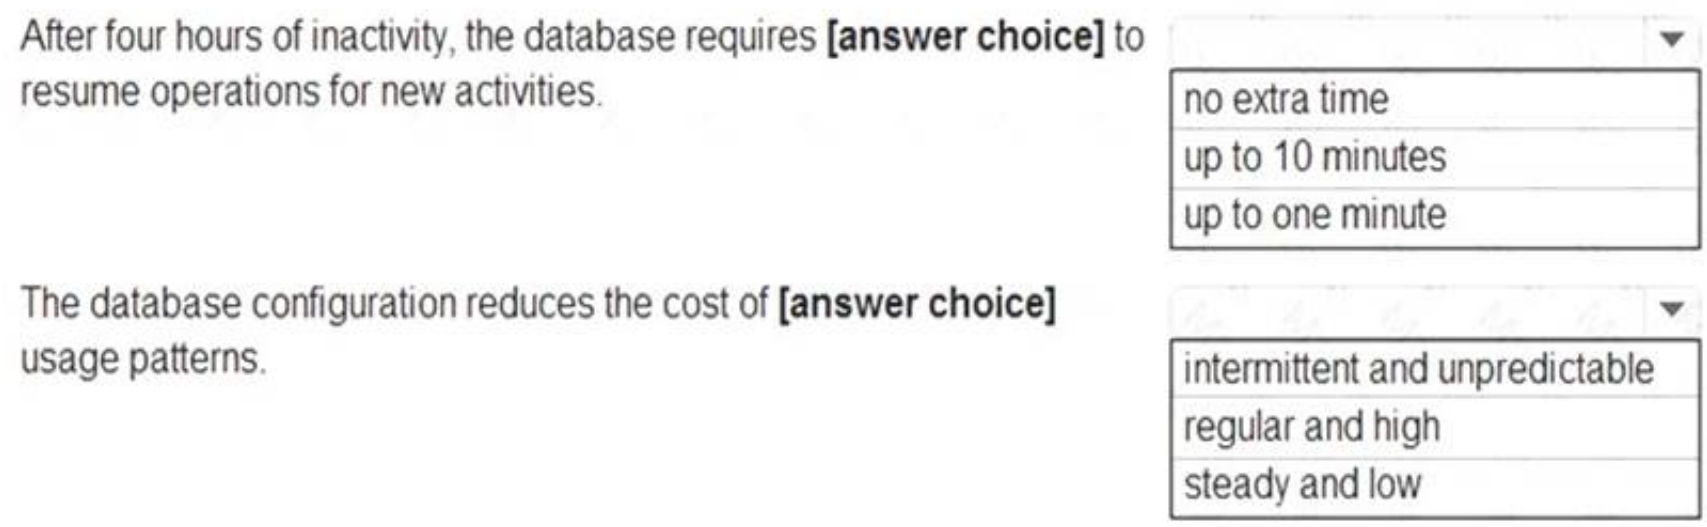

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview

NEW QUESTION 46
- (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 the dbcc dbreindex command.
C. Modify the automatic tuning settings for db1.
D. Modify the Query Store settings for db1.

Answer: C

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

NEW QUESTION 47
- (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 48
- (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 51
- (Exam Topic 5)
You have 50 Azure SQL databases.
You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.
What should you do?

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

Answer: D

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

NEW QUESTION 52
- (Exam Topic 5)

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

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

Answer: B

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

NEW QUESTION 55
- (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.

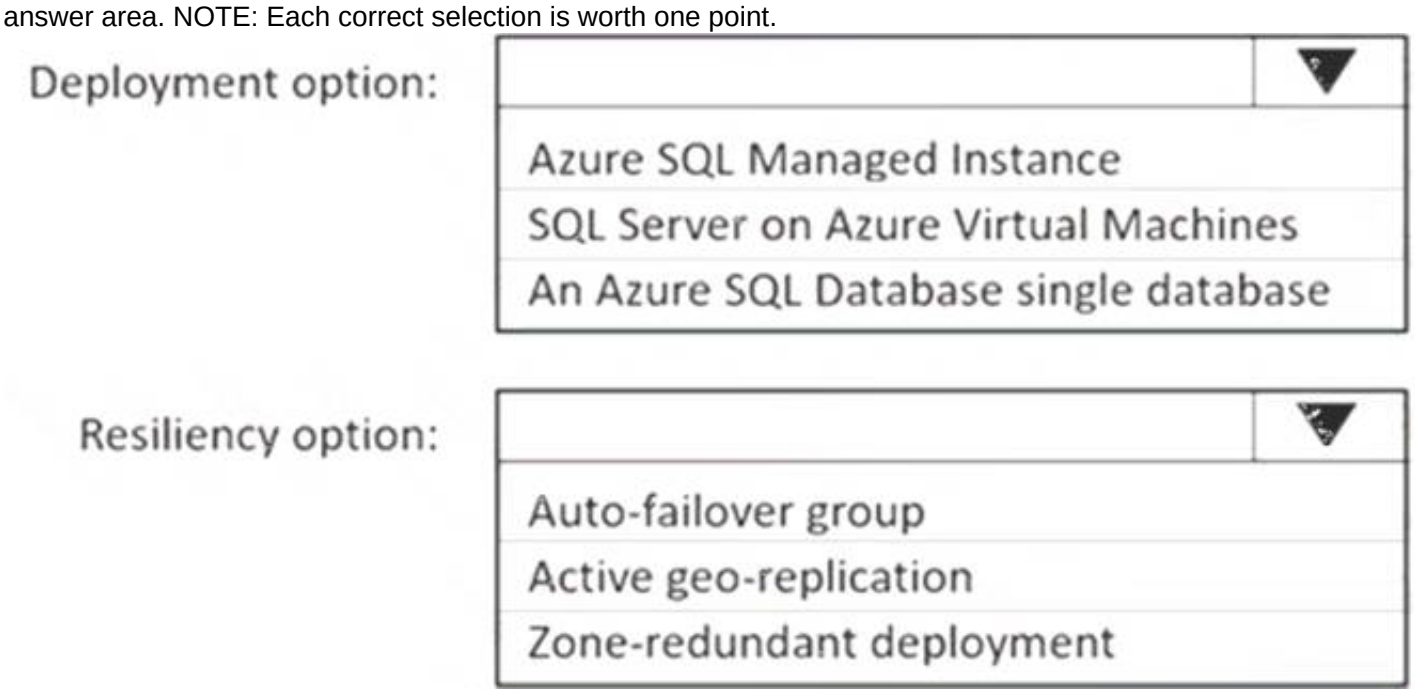

A. Mastered
B. Not Mastered

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 57
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine.
You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

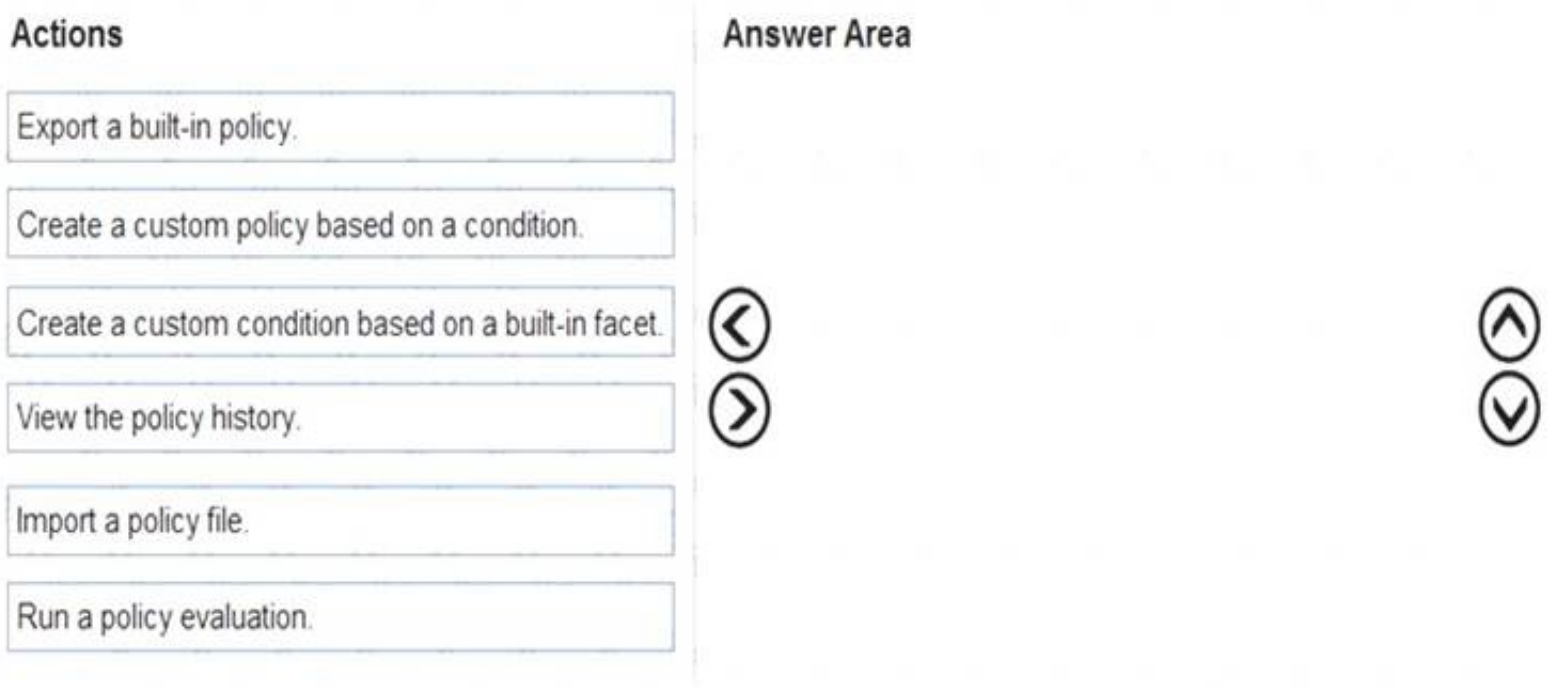

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 58
- (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.

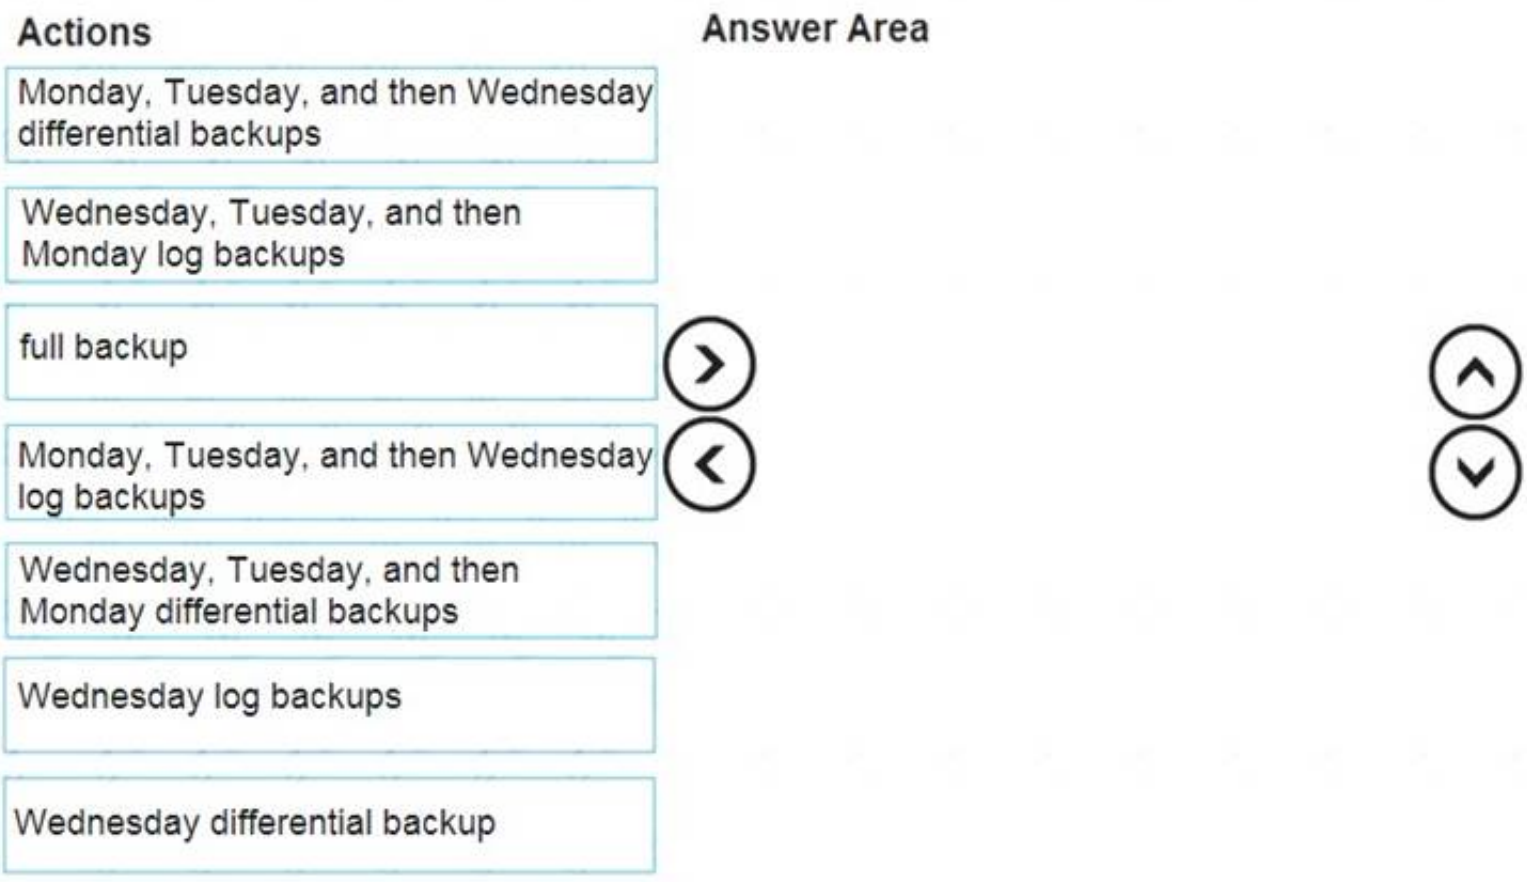

A. Mastered
B. Not Mastered

Answer: A

Explanation:

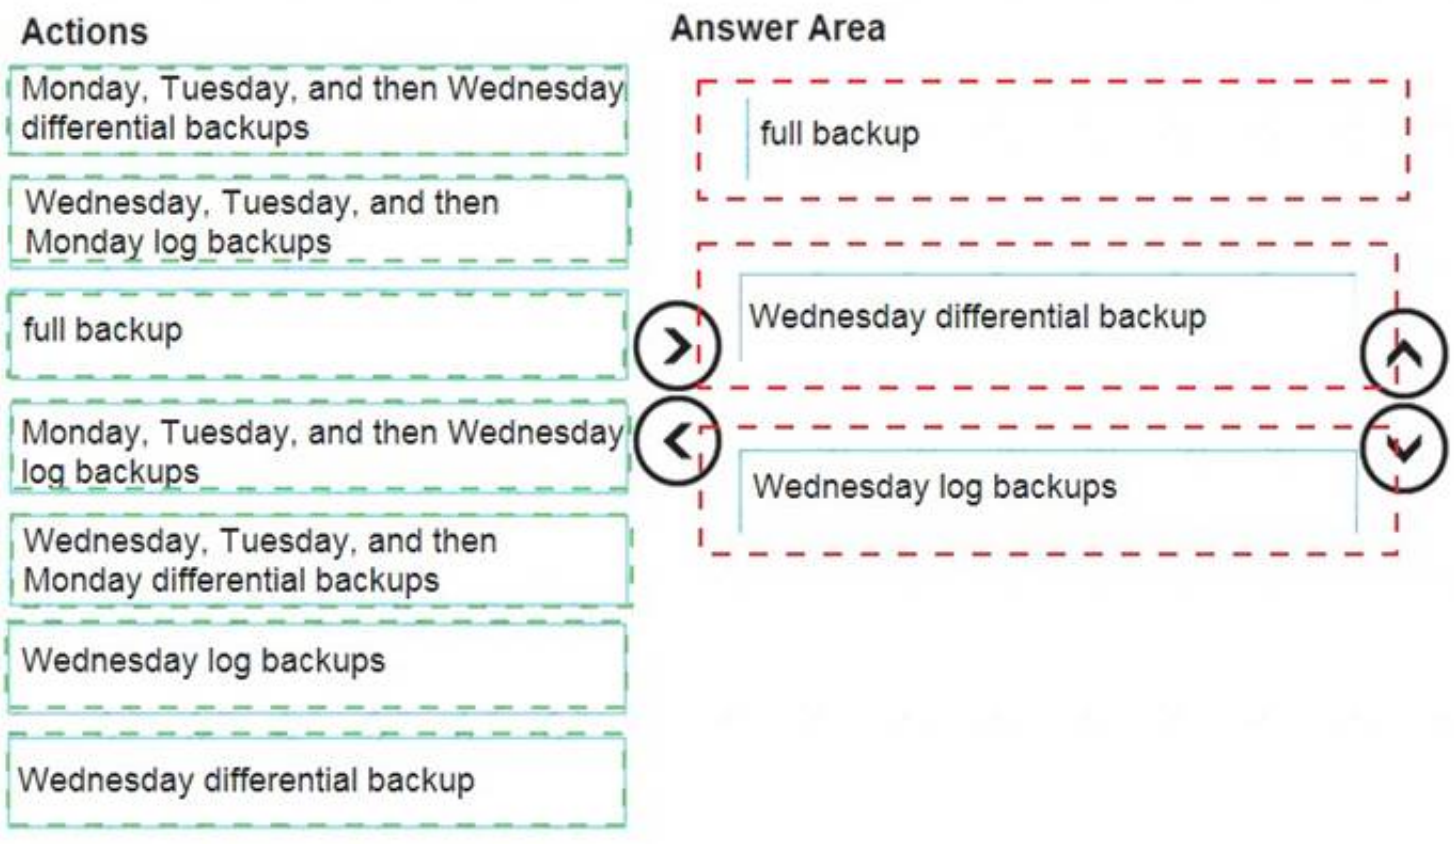

NEW QUESTION 59
- (Exam Topic 5)
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

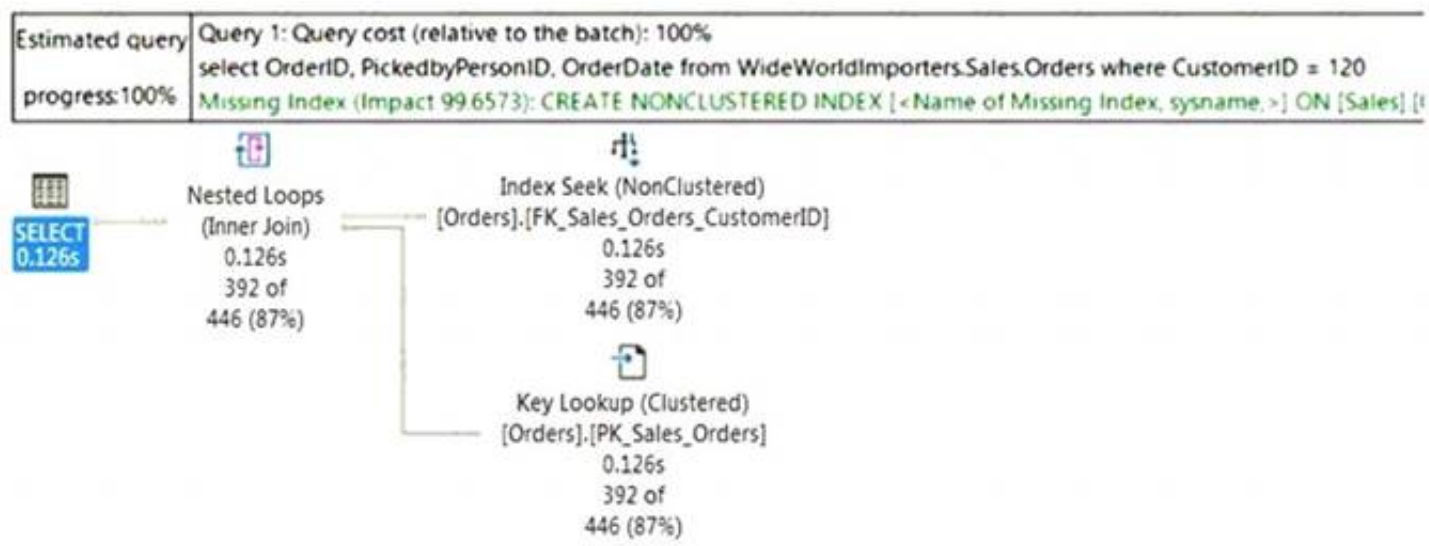

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

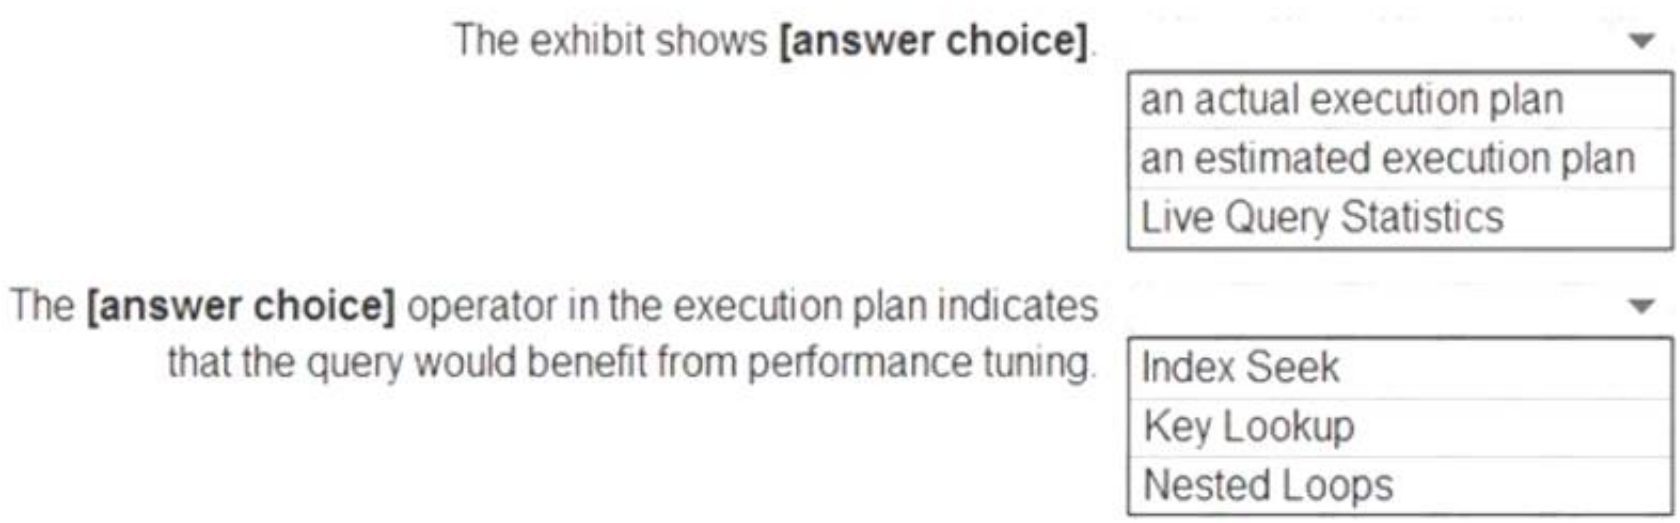

A. Mastered
B. Not Mastered

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 64
- (Exam Topic 5)
You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation.
Which service should you use to ingest the data?

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

Answer: D

Explanation:
The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.
Reference:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions

NEW QUESTION 66
- (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 67
- (Exam Topic 5)
You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:
- Send the output to an Azure Synapse.
- Identify spikes and dips in time series data.
- Minimize development and configuration effort. Which should you include in the solution?

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

Answer: C

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

NEW QUESTION 71
- (Exam Topic 5)
You have a on-premises Microsoft SQL Server named SQL1 that hosts five databases.
You need to migrate the databases to an Azure SQL managed instance. The solution must minimize downtime and prevent data loss.
What should you use?

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

Answer: A

NEW QUESTION 74
- (Exam Topic 5)
You have an Azure Stream Analytics job.
You need to ensure that the job has enough streaming units provisioned. You configure monitoring of the SU % Utilization metric.
Which two additional metrics should you monitor? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Late Input Events
B. Out of order Events

C. Backlogged Input Events
D. Watermark Delay
E. Function Events

Answer: CD

Explanation:
To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.
Note: 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 79
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.
Does this meet the goal?

A. Yes
B. No

Answer: A

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

NEW QUESTION 84
- (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 87
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.
You plan to migrate DB1 to an Azure SQL Database managed instance.
What should you use to minimize downtime and data loss during the migration?

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

Answer: D

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

NEW QUESTION 88
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machine named SQL1.
You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.
Where should you store the metrics?

A. a Log Analytics workspace
B. Azure Event Hubs
C. Azure SQL Database
D. an Azure Blob storage container

Answer: A

NEW QUESTION 92
- (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 95
- (Exam Topic 5)
You have a version-8.0 Azure Database for MySQL database.
You need to identify which database queries consume the most resources. Which tool should you use?

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

Answer: A

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

NEW QUESTION 100
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine named SQL1. SQL1 has an agent job to back up all databases.
You add a user named dbadmin1 as a SQL Server Agent operator. You need to ensure that dbadmin1 receives an email alert if a job fails.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

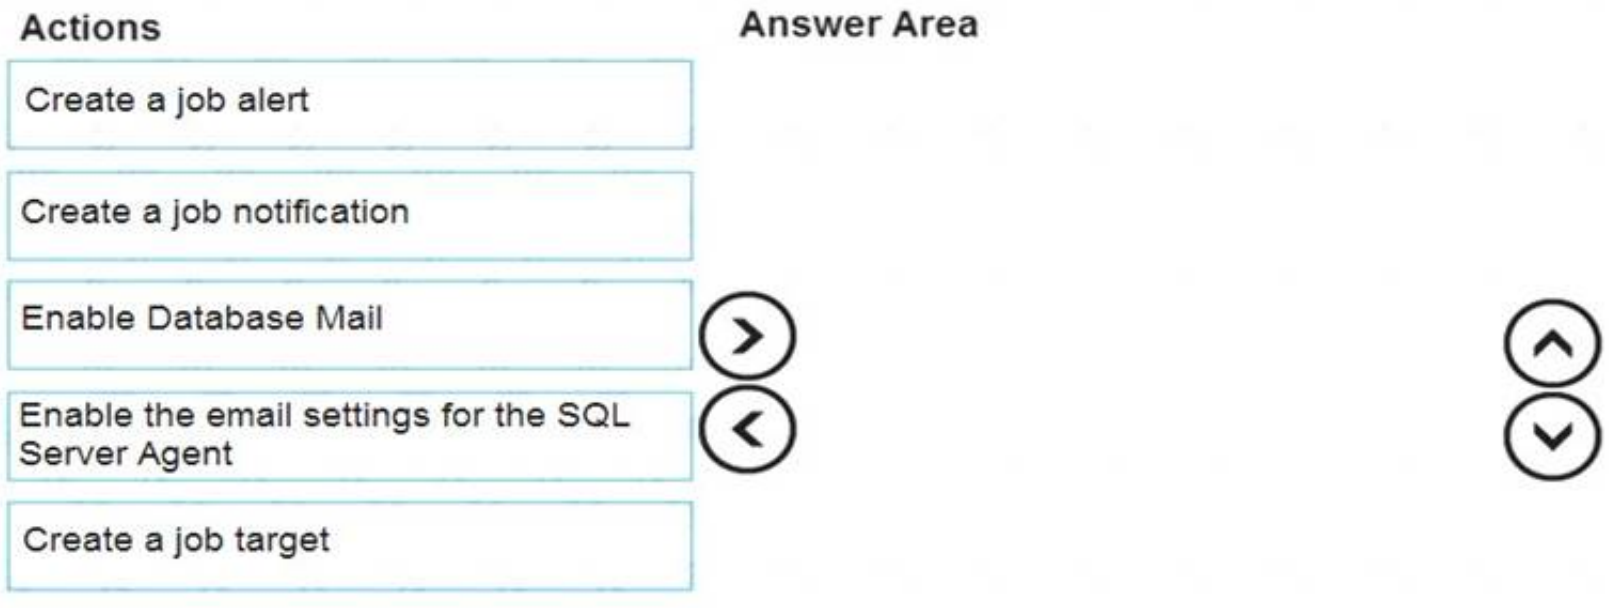

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Enable the email settings for the SQL Server Agent.
To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.
Step 2: Create a job alert
Step 3: Create a job notification Example:
-- adds an e-mail notification for the specified alert (Test Alert)
-- This example assumes that Test Alert already exists
-- and that François Ajenstat is a valid operator name. USE msdb ;
GO
EXEC dbo.sp_add_notification
@alert_name = N'Test Alert',
@operator_name = N'François Ajenstat',
@notification_method = 1 ; GO
Reference:
https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator

NEW QUESTION 103
- (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:
Text Description automatically generated with low confidence
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-serv

NEW QUESTION 106
- (Exam Topic 5)
You create five Azure SQL Database instances on the same logical server.

In each database, you create a user for an Azure Active Directory (Azure AD) user named User1. User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.
You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.
What should you do?

A. Create User1 in the master database.
B. Assign User1 the db_datareader role for the master database.
C. Assign User1 the db_datareader role for the databases that Userl creates.
D. Grant select on sys.databases to public in the master database.

Answer: A

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

NEW QUESTION 108
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:
- During normal operations, provide at least two readable copies of Sales.
- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication. Does this meet the goal?

A. Yes
B. No

Answer: B

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

NEW QUESTION 110
- (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 114
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named Db1.
You need to enable automatic tuning for Db1.
How should you complete the statements? To answer, select the appropriate answer in the answer area.
NOTE: Each correct selection is worth one point.

ALTER DATABASE [Db1]

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

GO
ALTER DATABASE [Db1]

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

GO

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SET AUTOMATIC_TUNING = AUTO
To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:
ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO
Setting automatic tuning to AUTO will apply Azure Defaults.
Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)
To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:
ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)
Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option. Setting it to OFF will also override any setting that database inherited and disable the tuning option.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable

NEW QUESTION 118
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database. The database fails to respond to queries in a timely manner.
You need to identify whether the issue relates to resource_semaphore waits.
How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
SELECT
```

is_user_process
wait_time
wait_type

```
   SUM(wait_time) AS total_wait_time_ms
FROM sys.
```

dm_exec_query_stats
dm_exec_requests
query_store_query

```
  JOIN sys.dm_exec_sessions AS dmv2
     ON dmv1.session_id = dmv2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs

NEW QUESTION 123
- (Exam Topic 5)
You receive numerous alerts from Azure Monitor for an Azure SQL database.
You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Answer: BD

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

NEW QUESTION 124
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to display the estimated execution plan of a query by using the query editor in the Azure portal. What should you do first?

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

Answer: A

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

NEW QUESTION 125
- (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.

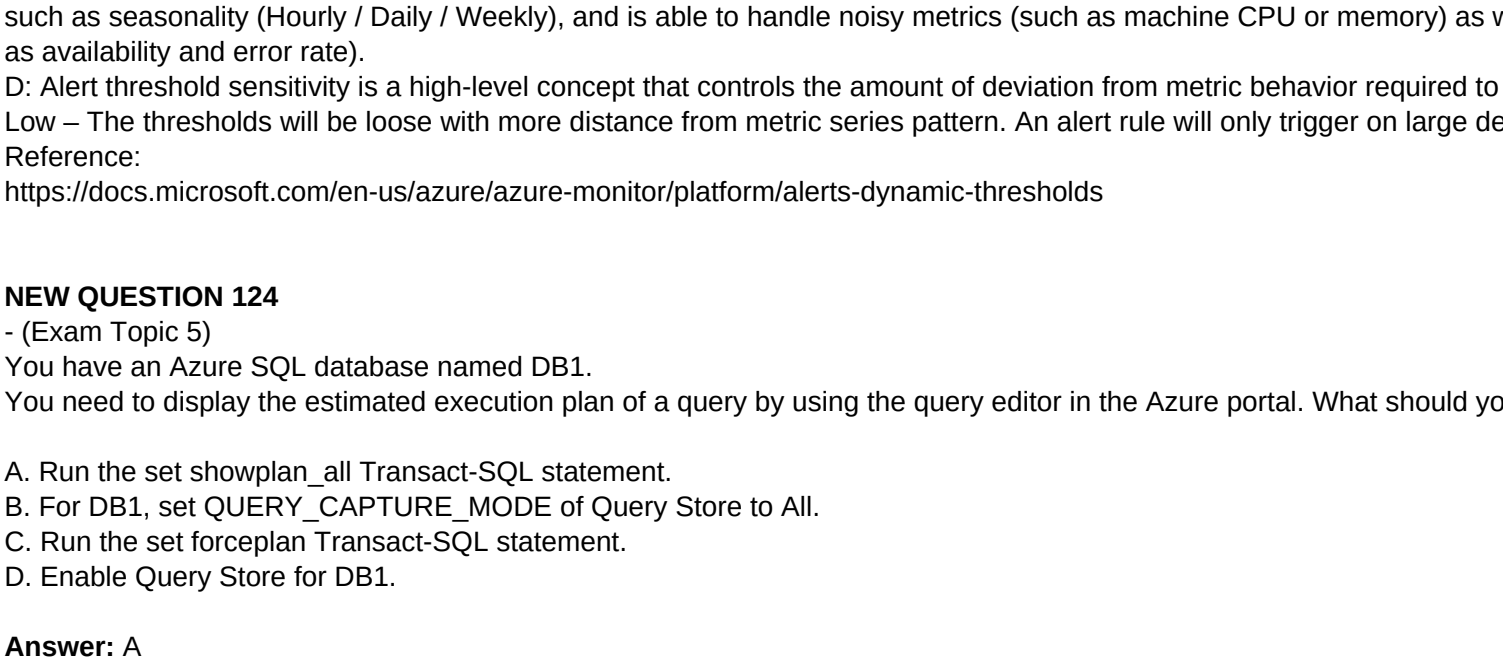

A. Mastered
B. Not Mastered

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 127
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. 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 132
- (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: A

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 136
- (Exam Topic 5)
You have an Azure subscription that contains a logical SQL server named Serverl. The master database of Serverl contains a user named Userl. You need to ensure that Userl can create databases on Serverl. Which database role should you assign to Userl?

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

Answer: B

NEW QUESTION 138
- (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 143
- (Exam Topic 4)
You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements.
What should you create?

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

Answer: B

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

NEW QUESTION 144
- (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

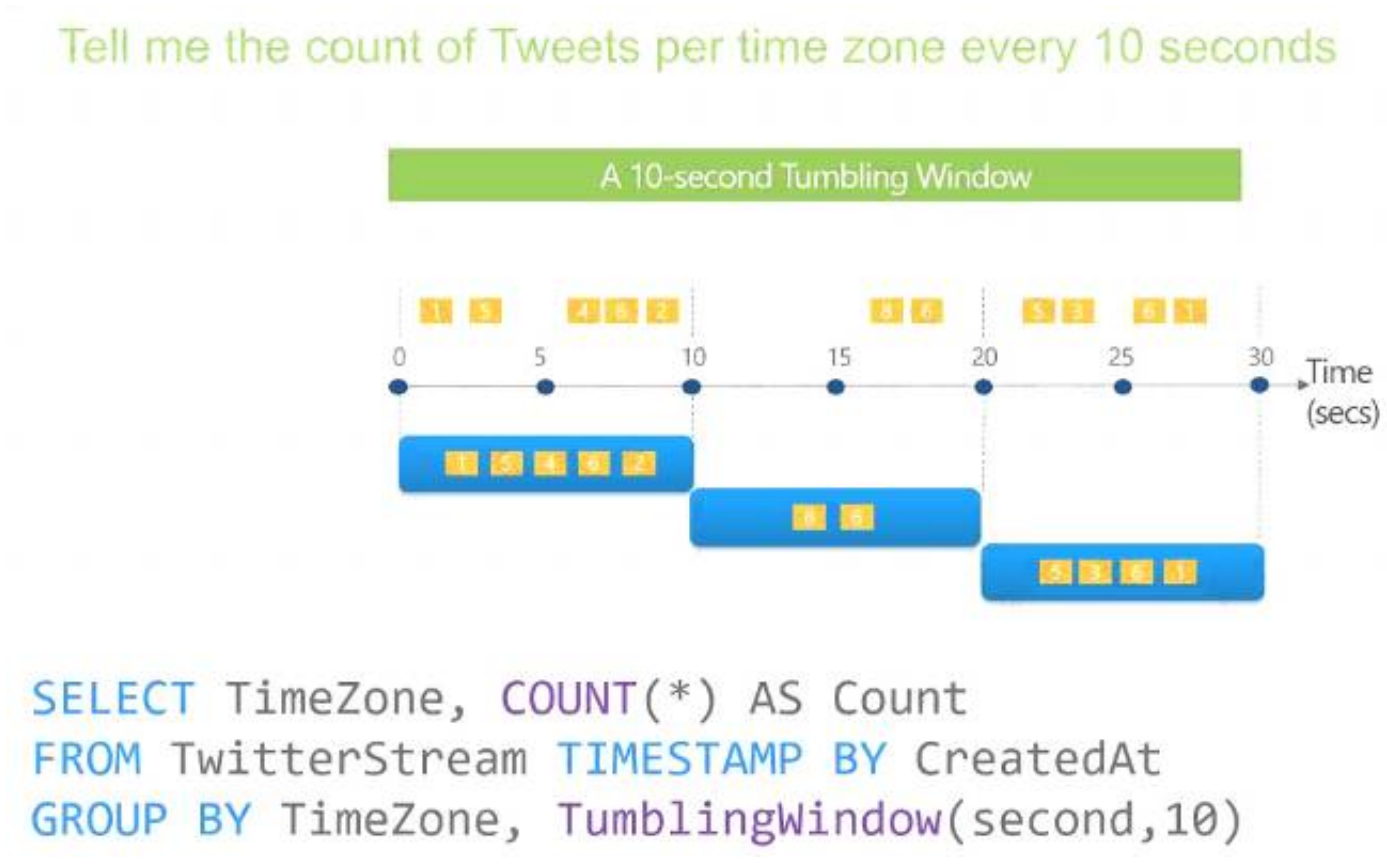

Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 145
- (Exam Topic 2)
What should you use to migrate the PostgreSQL database?

A. Azure Data Box
B. AzCopy
C. Azure Database Migration Service
D. Azure Site Recovery

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/dms/dms-overview

NEW QUESTION 150
- (Exam Topic 1)
You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals.
What should you include in the recommendation?

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

Answer: D

NEW QUESTION 155
- (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.

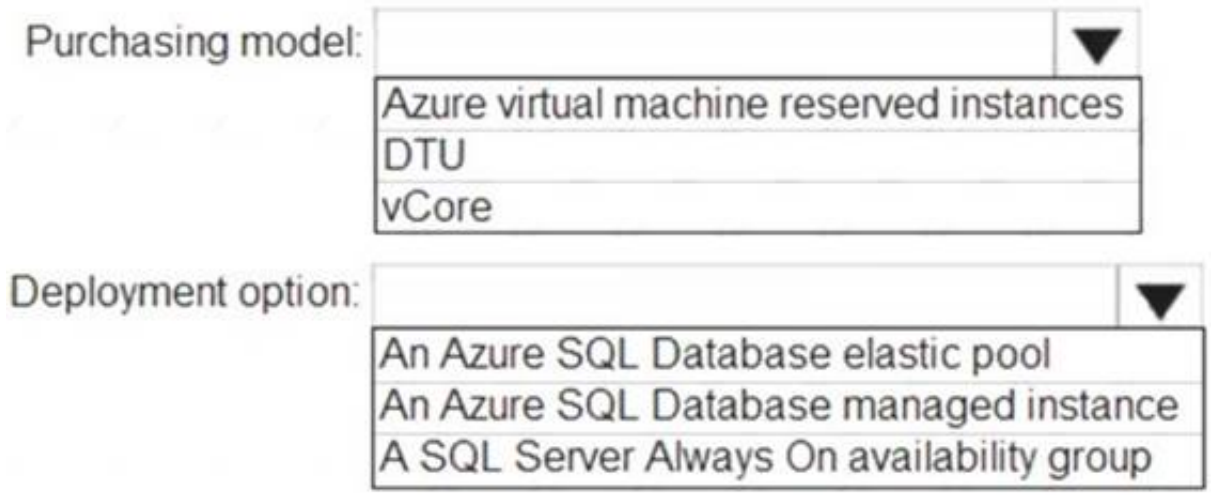

A. Mastered
B. Not Mastered

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 160
.....

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