Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)

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

NEW QUESTION 1
- (Exam Topic 5)
You have an Azure SQL database.
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 2
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named
container1.
You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will
produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column. 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.
Note: In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

NEW QUESTION 3
- (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 4
- (Exam Topic 5)
Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.
You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

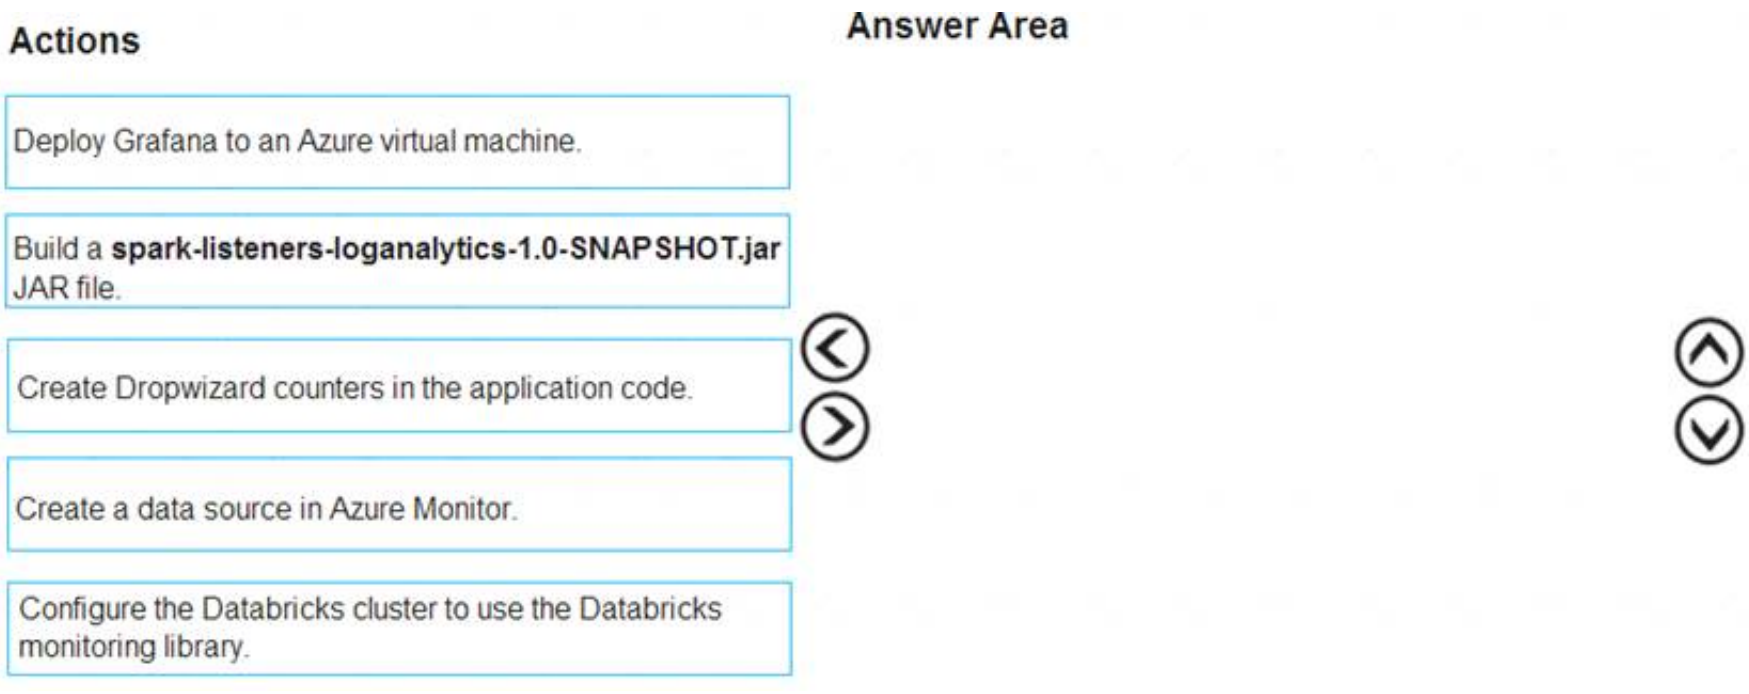

A. Mastered
B. Not Mastered

Answer: A

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

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

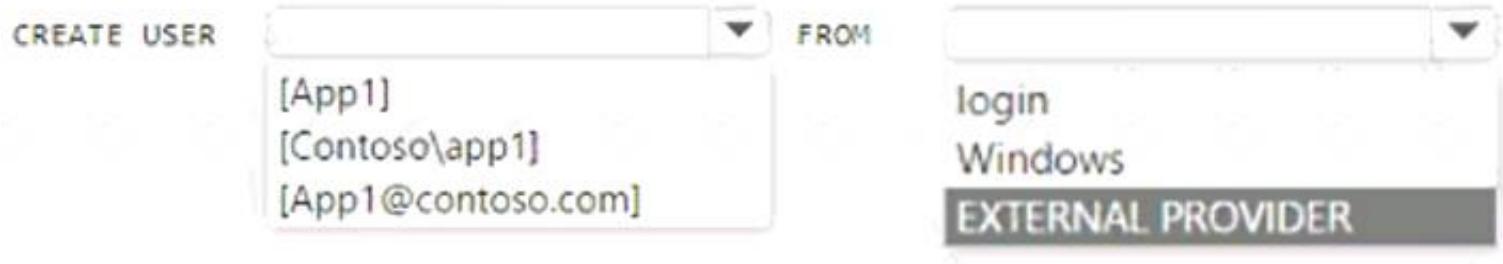

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 6
- (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 7
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

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

NEW QUESTION 8
- (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 9
- (Exam Topic 5)
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.
You need to ensure that User1 can perform the following tasks on DB1:
• Create jobs.
• View all jobs.
• Modify, delete, and disable the jobs the user created. The solution must use the principle of least privilege.
Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

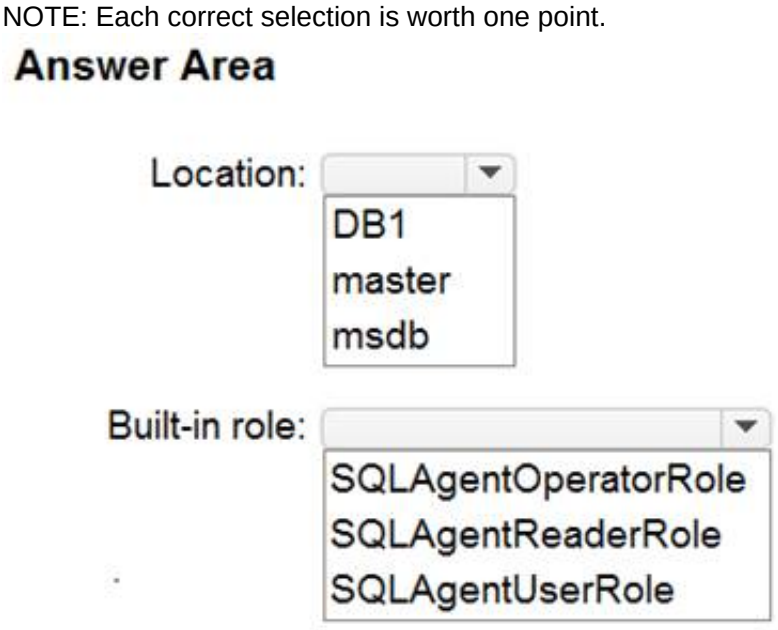

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 10

- (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 10
- (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 13
- (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 18
- (Exam Topic 5)
You manage 100 Azure SQL managed instances located across 10 Azure regions.
You need to receive voice message notifications when a maintenance event affects any of the 10 regions. The solution must minimize administrative effort.
What should you do?

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

Answer: C

NEW QUESTION 20
- (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.

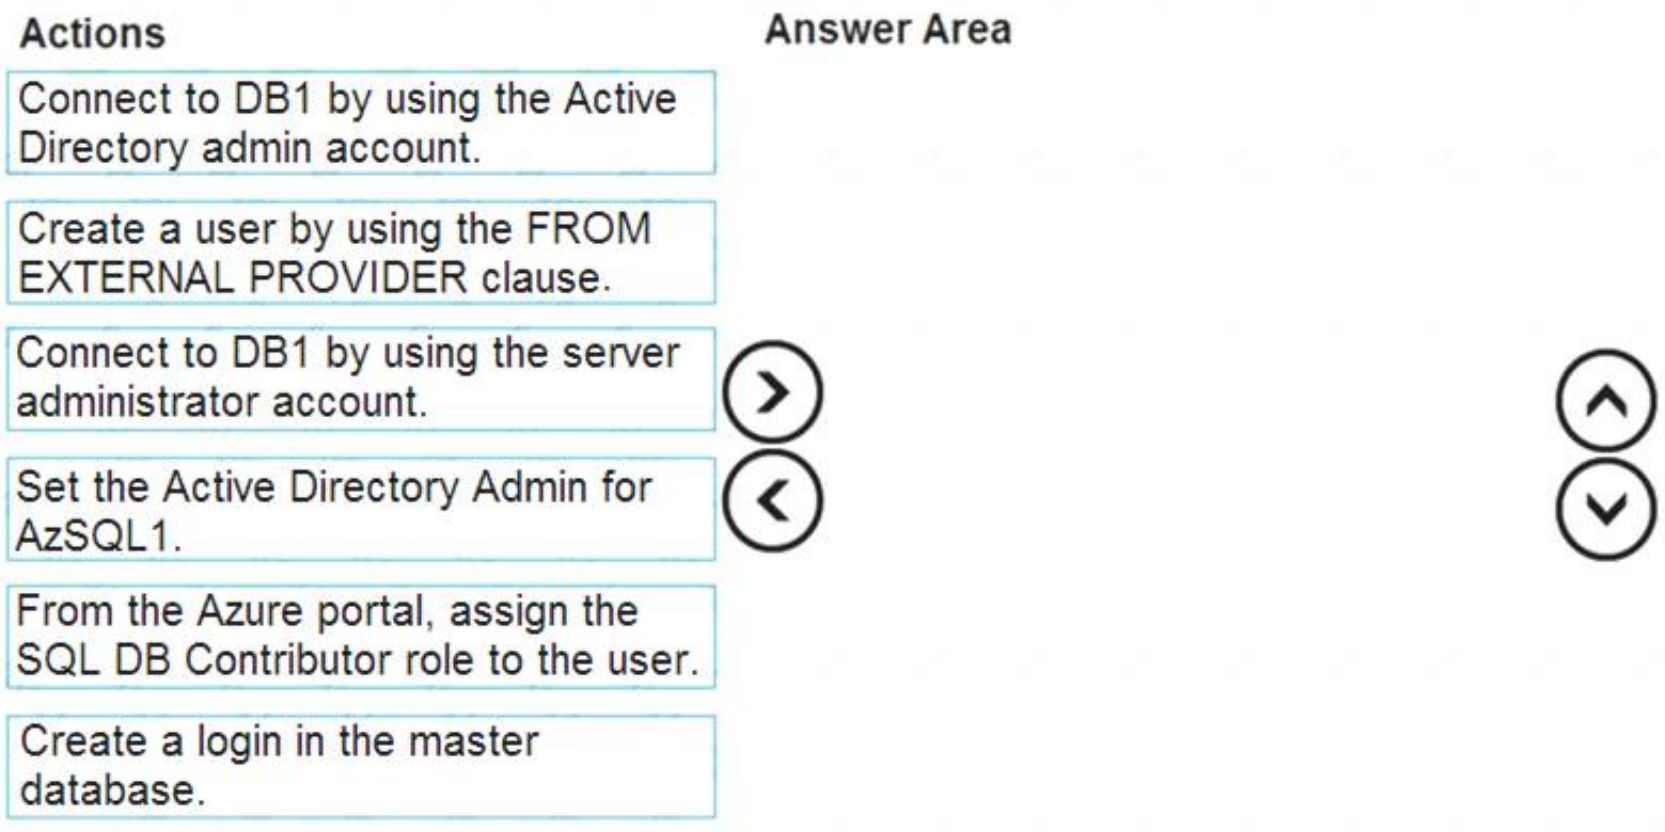

A. Mastered
B. Not Mastered

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 25
- (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.

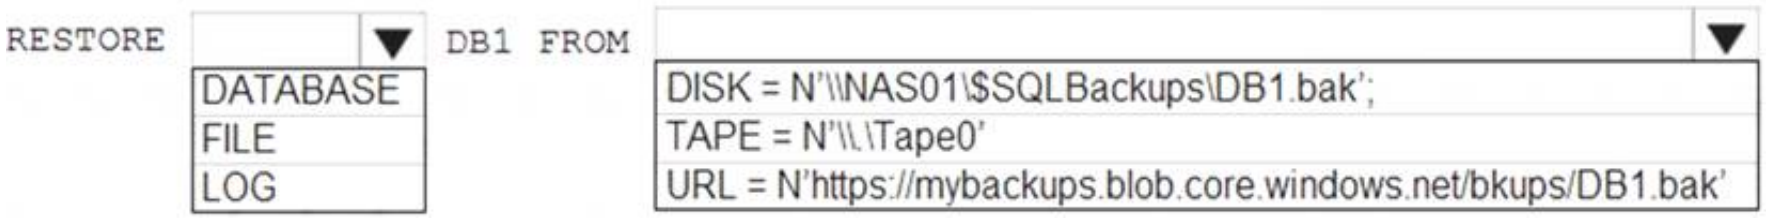

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated

NEW QUESTION 27
- (Exam Topic 5)
You have an Azure SQL database named DB1 in the General Purpose service tier. You need to monitor D81 by using SQL insights.
What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

To collect monitoring data, use:
- A virtual machine
- An Azure function
- The Azure Monitor agent

To store monitoring data, create:
- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

To collect monitoring data, use:
- A virtual machine
- An Azure function
- The Azure Monitor agent

To store monitoring data, create:
- A Log Analytics workspace
- An Azure SQL database
- An Azure Storage account

NEW QUESTION 32
- (Exam Topic 5)
You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases. You need to update the column and index statistics for the databases.
What should you use?

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

Answer: A

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

NEW QUESTION 35
- (Exam Topic 5)
You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs
Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

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

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:

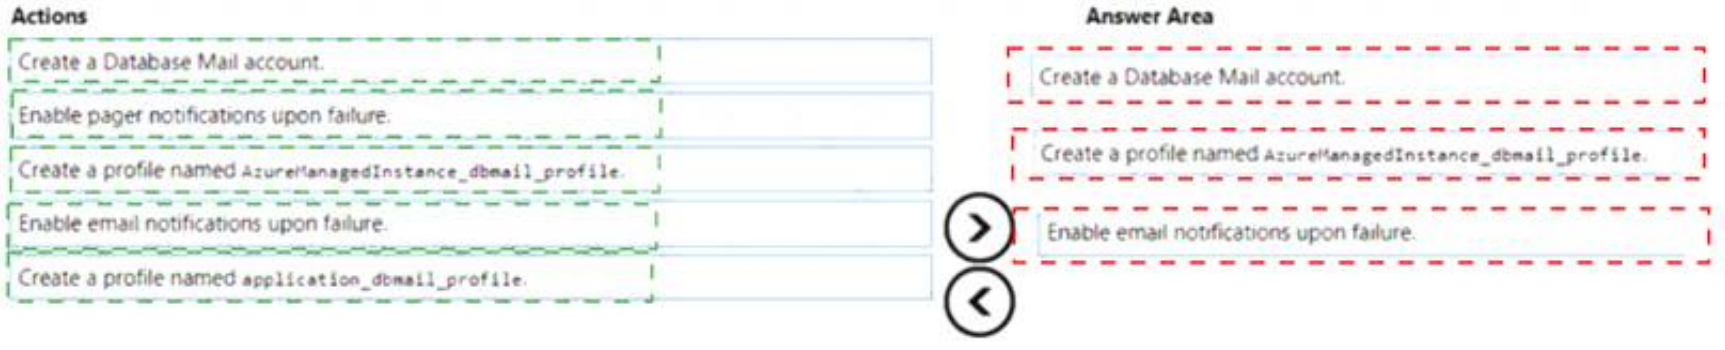

NEW QUESTION 38
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct 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 39
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

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

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

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

Answer: D

NEW QUESTION 41
- (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 43
- (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 47
- (Exam Topic 5)
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

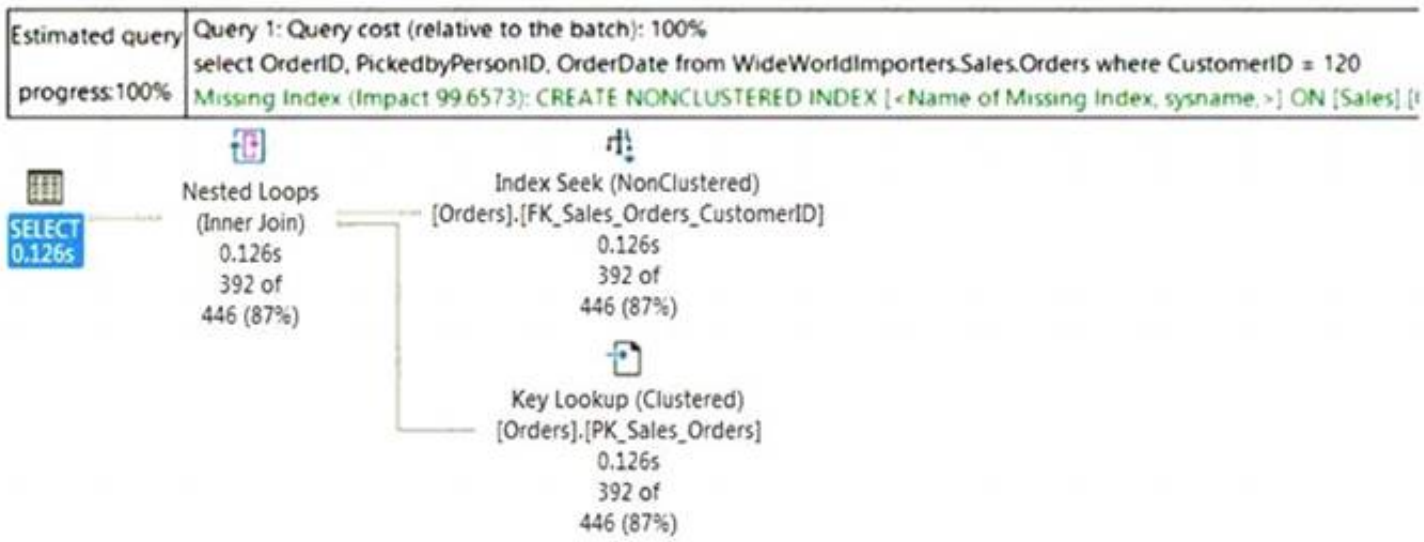

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE:Each correct selection is worth one point.

The exhibit shows **[answer choice]**.
- an actual execution plan
- an estimated execution plan
- Live Query Statistics

The **[answer choice]** operator in the execution plan indicates that the query would benefit from performance tuning.
- Index Seek
- Key Lookup
- Nested Loops

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 51
- (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.

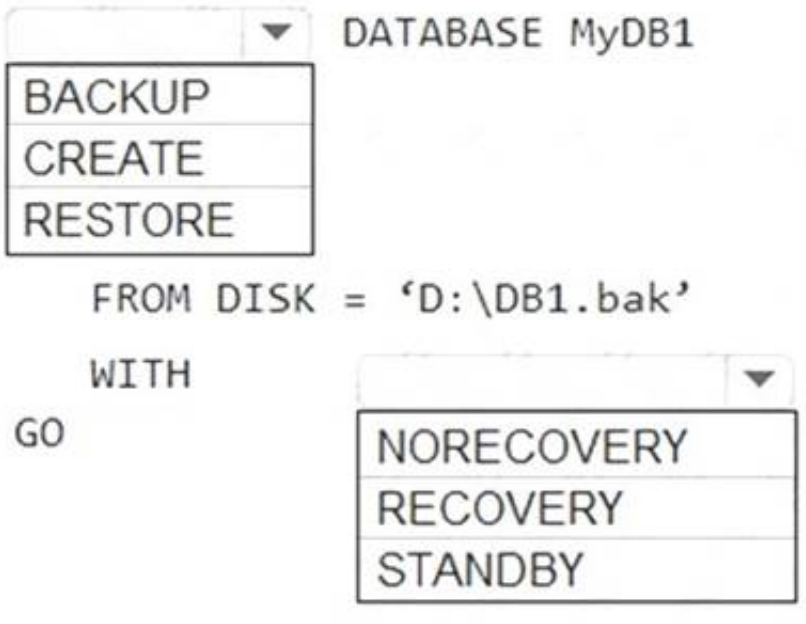

A. Mastered
B. Not Mastered

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 53
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 contains a table named CustomerPII.
You need to record whenever users query the CustomerPII table.
Which two options should you enable? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: AC

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

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

Note:
The Server Audit Specification object belongs to an audit.
A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auditing-overview

NEW QUESTION 54
- (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 59
- (Exam Topic 5)
You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD sen/ice principal named appl.
You need to ensure that appl can access db1. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the
answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

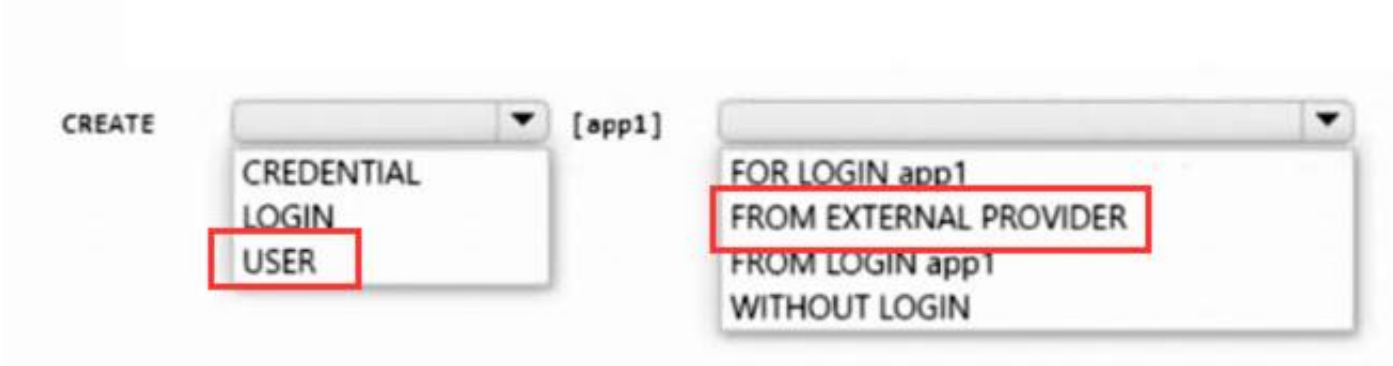

NEW QUESTION 61
- (Exam Topic 5)
You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.
You need to ensure that the design provides the fastest query times of the records when querying for arbitrary 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 65
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. You have an application that queries DB1 to generate a sales report.
You need to see the parameter values from the last time the query was executed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: AC

Explanation:
Last_Query_Plan_Stats allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in sys.dm_exec_query_plan_stats.
Lightweight profiling can be disabled at the database level using the LIGHTWEIGHT_QUERY_PROFILING database scoped configuration: ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure

NEW QUESTION 67
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.
You need to resolve the performance issues. What should you do?

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

Answer: A

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

NEW QUESTION 71
- (Exam Topic 5)
You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.
You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.
Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

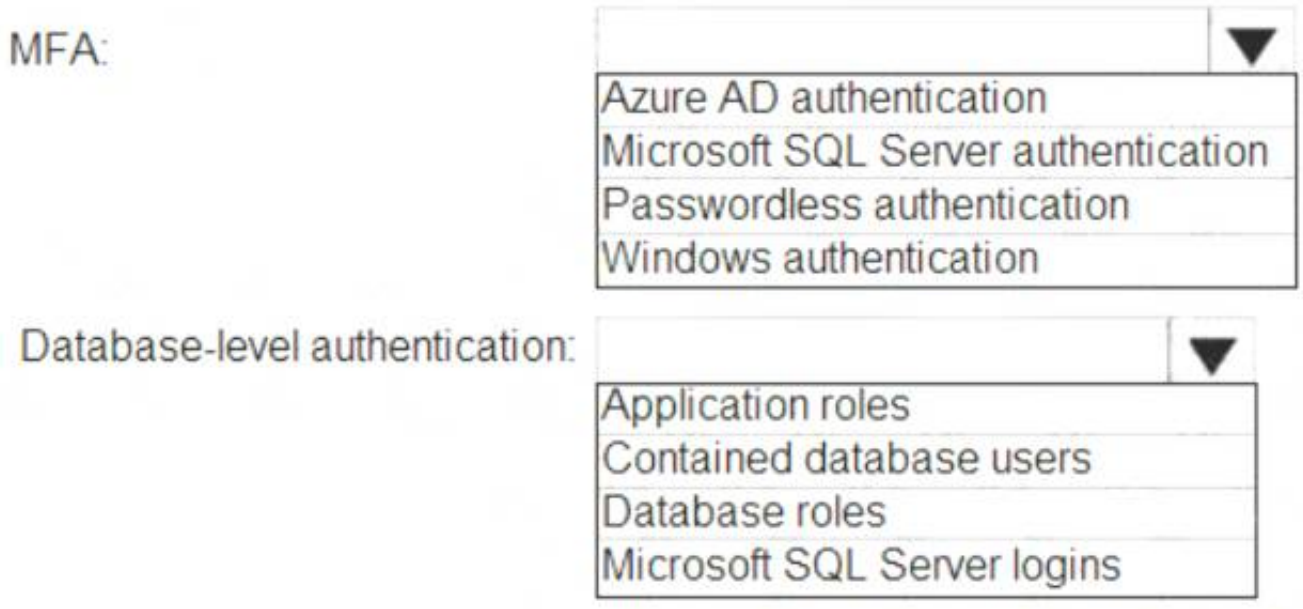

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 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 75
- (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 76
- (Exam Topic 5)
You have an Azure SQL database.
You need to identify whether a delayed query execution is associated to a RESOURCE wait.
How should you complete the Transact –SQL statement? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

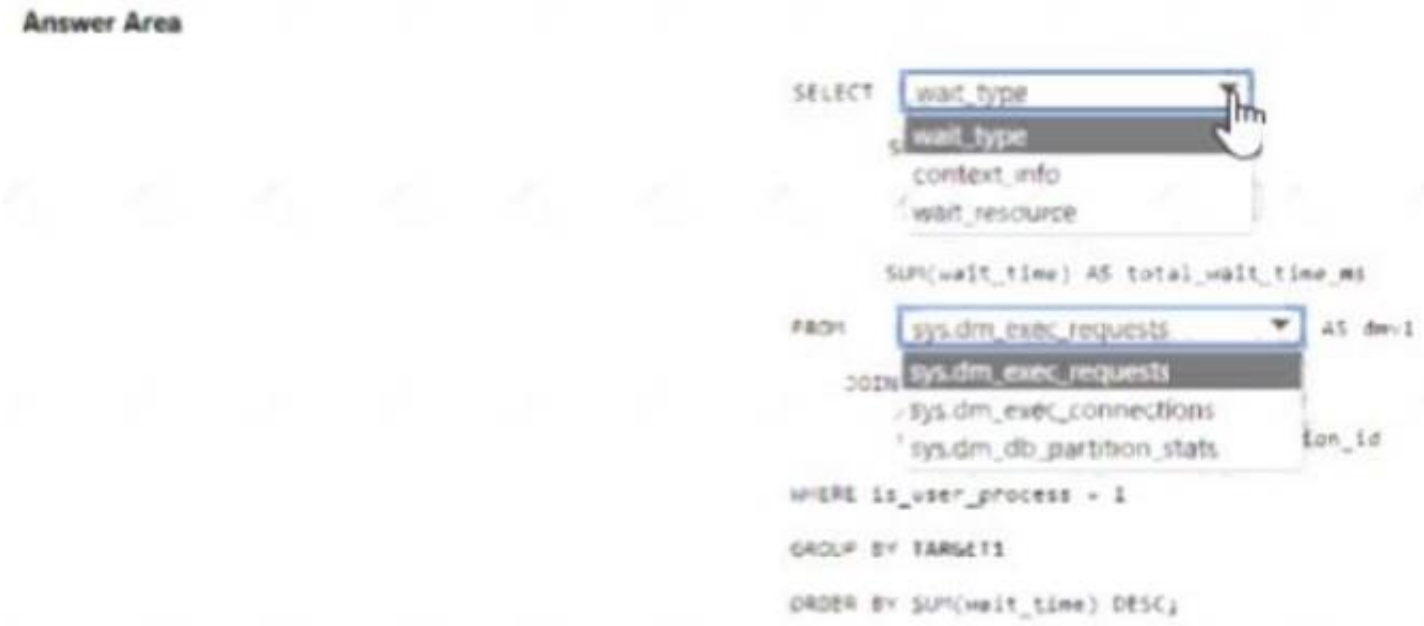

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

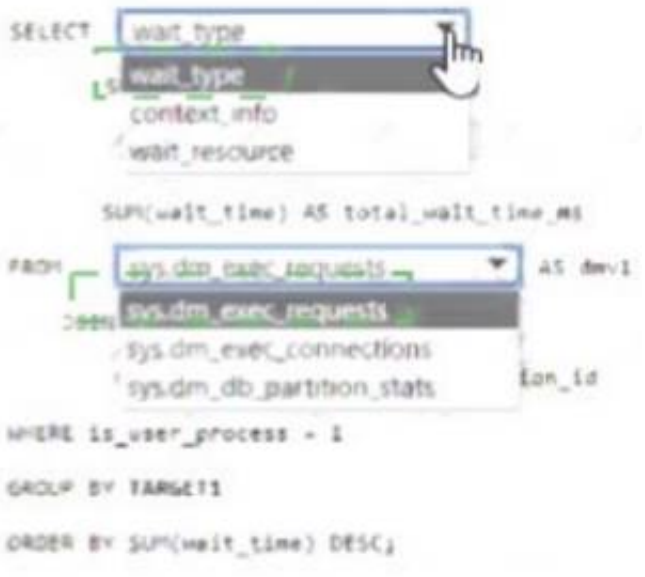

NEW QUESTION 80
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication. You need to migrate SQL1 to Azure. Which service should you use?

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

Answer: C

NEW QUESTION 85
- (Exam Topic 5)
You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features. You plan to migrate to Azure SQL.
Which service should you use?

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

Answer: B

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

NEW QUESTION 88
- (Exam Topic 5)
You have an Azure SQL Database elastic pool that contains 10 databases. You receive the following alert.
Msg 1132, Level 16, State 1, Line 1
The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.
You need to resolve the alert. The solution must minimize administrative effort.
Which three actions can you perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

Answer: BCD

NEW QUESTION 91
- (Exam Topic 5)
You configure backups for an Azure SQL database as shown in the following exhibit.

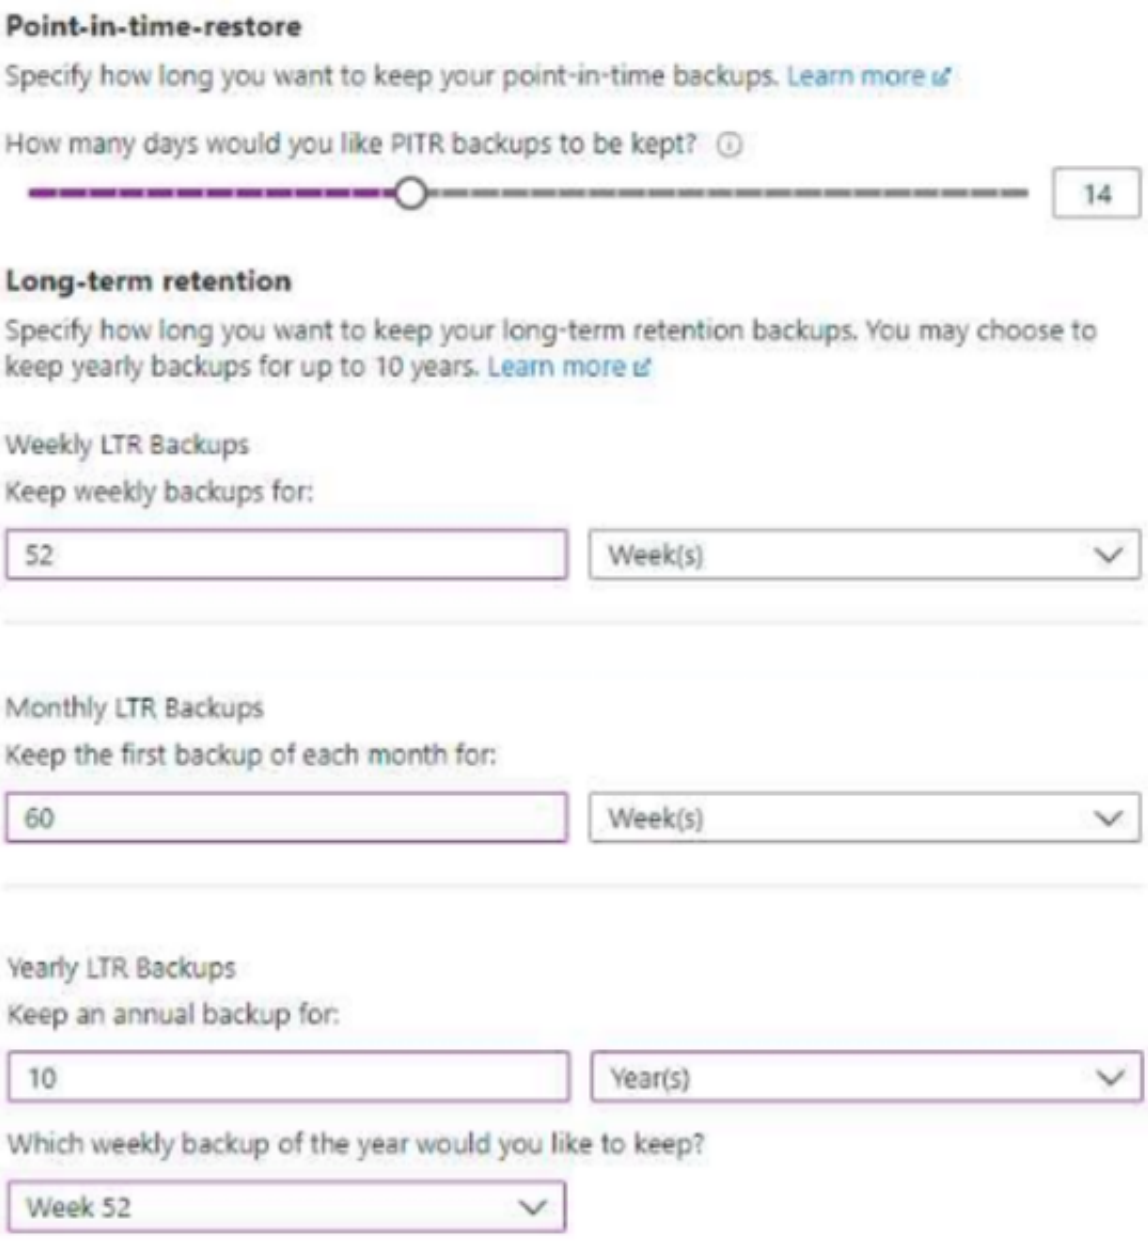

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

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 96
- (Exam Topic 5)
You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created m 06I The solution must use the principle of least privilege Howshould 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

GRANT [SELECT / DELETE / INSERT / SELECT / UPDATE] [Schema / Database / Schema / Table] ::[table] TO [Contoso\group1]

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

GRANT [SELECT / DELETE / INSERT / SELECT / UPDATE] [Schema / Database / Schema / Table] ::[table] TO [Contoso\group1]

NEW QUESTION 99
- (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 102
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica. What should you do? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Answer Area

To add read-only replicas of DB1:
- Create a replica on the same logical server.
- Create a new logical server and configure geo-replication.
- Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:
- Add an ApplicationIntent entry to the connection string.
- Add a MultiSubnetFailover entry to the App1 connection string.
- Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://sqlserverguides.com/read-only-replica-azure-sql/

NEW QUESTION 107
- (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 108
- (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 113
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

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

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

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

Answer: D

NEW QUESTION 116
- (Exam Topic 5)
You are building an Azure Stream Analytics job to retrieve game data.
You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.
How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

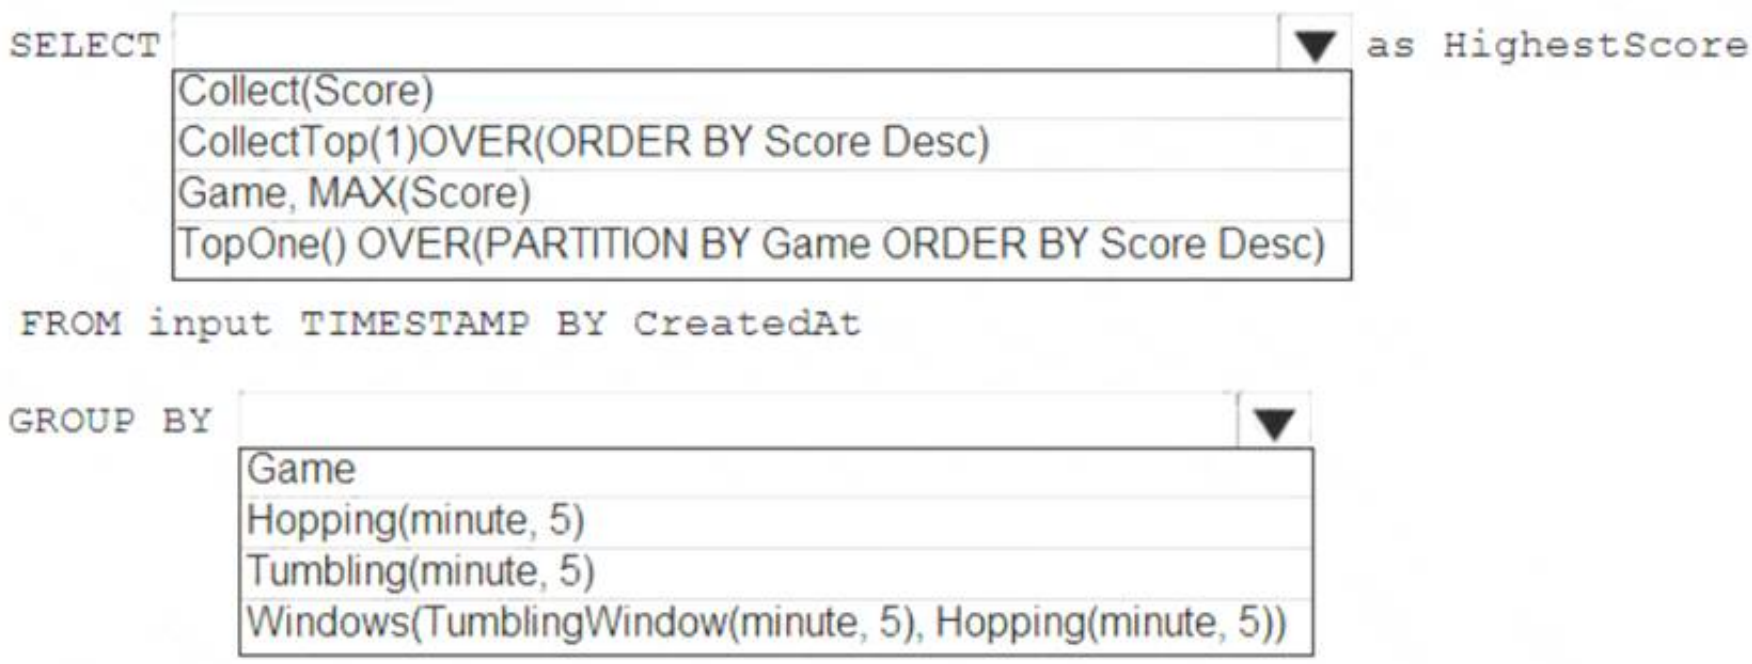

A. Mastered
B. Not Mastered

Answer: A

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

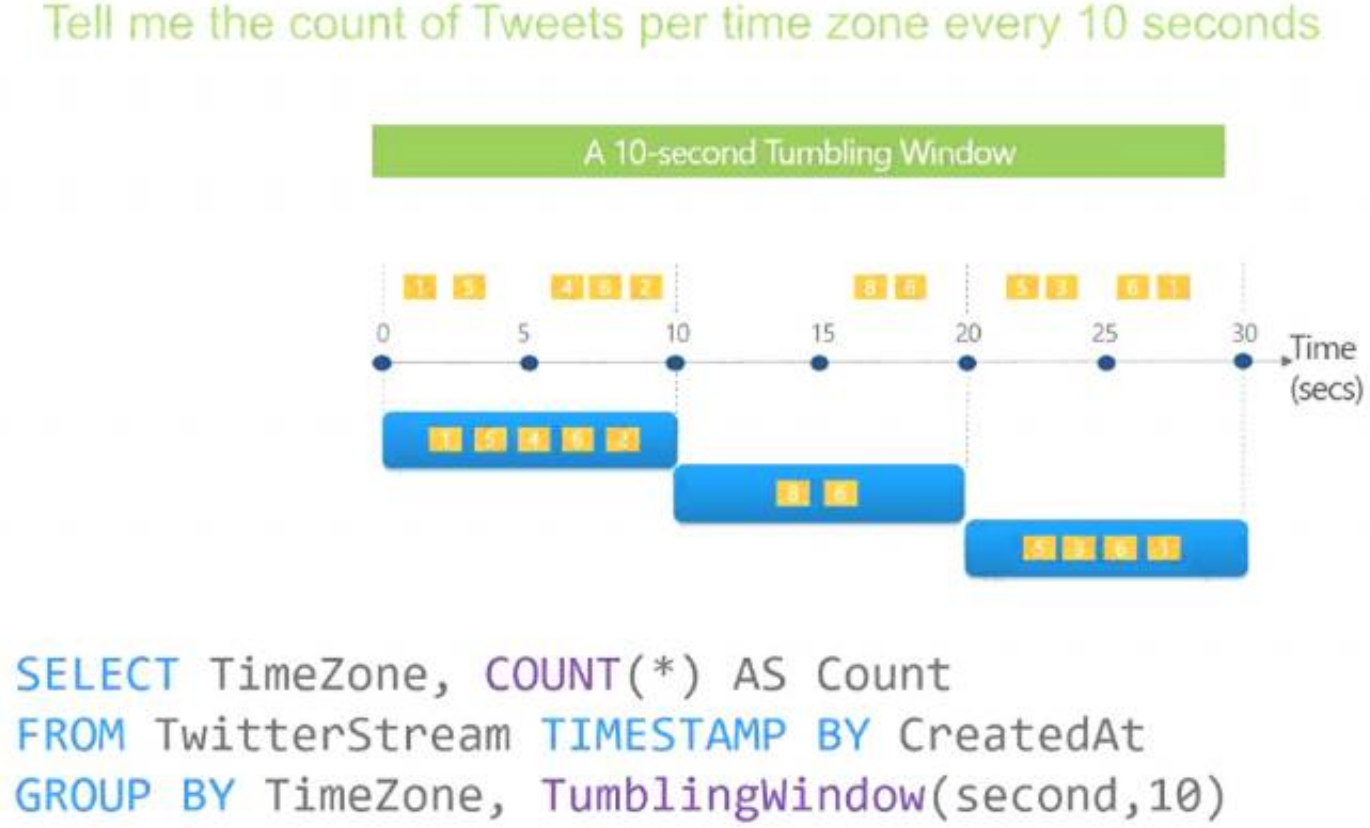

Reference:
https://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 121
- (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

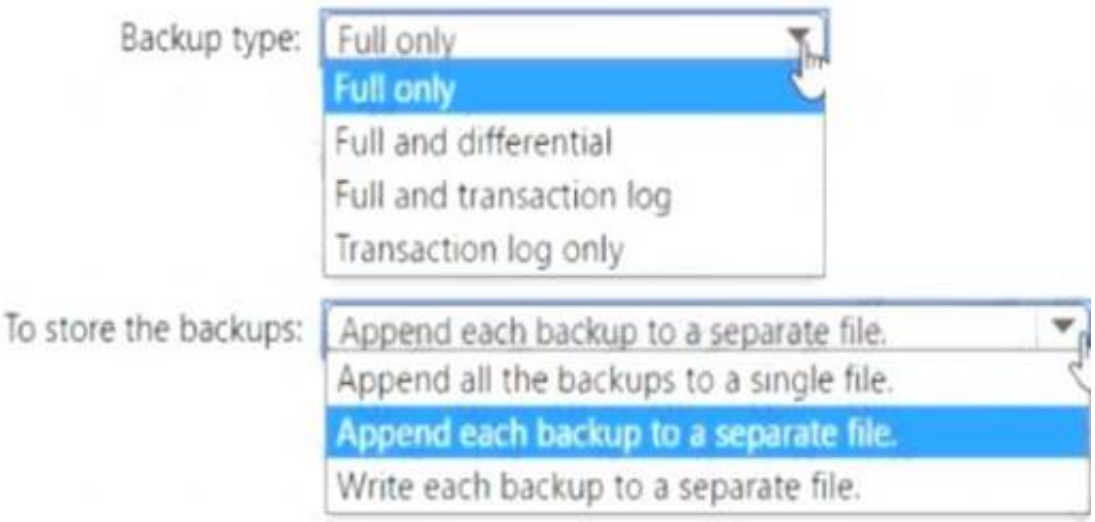

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Answer Area

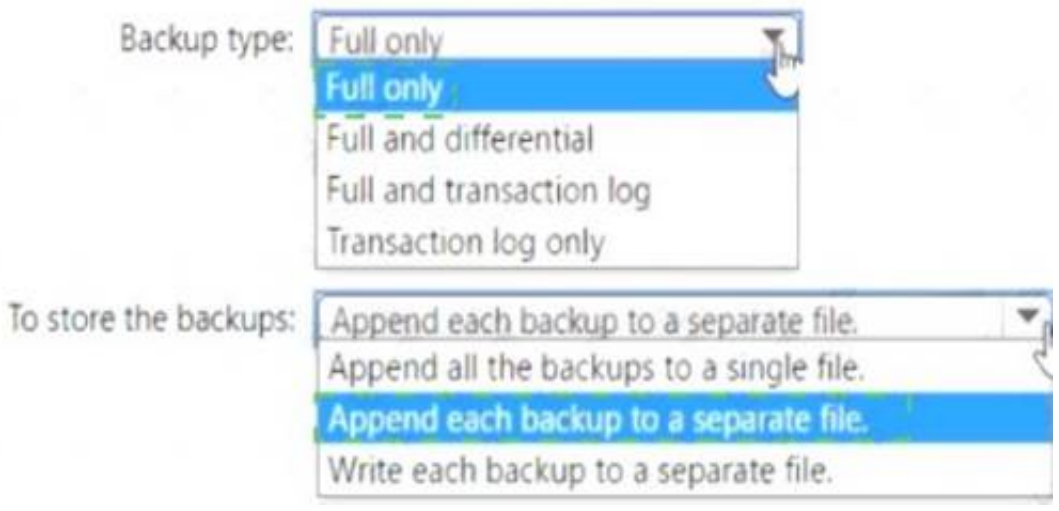

NEW QUESTION 124
- (Exam Topic 5)
A company plans to use Apache Spark analytics to analyze intrusion detection data.
You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts.
What should you recommend?

A. Azure Data Lake Storage
B. Azure Databricks
C. Azure HDInsight
D. Azure Data Factory

Answer: C

Explanation:
Azure HDInsight offers pre-made, monitoring dashboards in the form of solutions that can be used to monitor the workloads running on your clusters. There are solutions for Apache Spark, Hadoop, Apache Kafka, live long and process (LLAP), Apache HBase, and Apache Storm available in the Azure Marketplace.
Note: With Azure HDInsight you can set up Azure Monitor alerts that will trigger when the value of a metric or the results of a query meet certain conditions. You can condition on a query returning a record with a value that is greater than or less than a certain threshold, or even on the number of results returned by a query. For example, you could create an alert to send an email if a Spark job fails or if a Kafka disk usage becomes over 90 percent full.
Reference:
https://azure.microsoft.com/en-us/blog/monitoring-on-azure-hdinsight-part-4-workload-metrics-and-logs/

NEW QUESTION 126
- (Exam Topic 5)
You configure backup for an Azure SQL database as shown in the following exhibit.

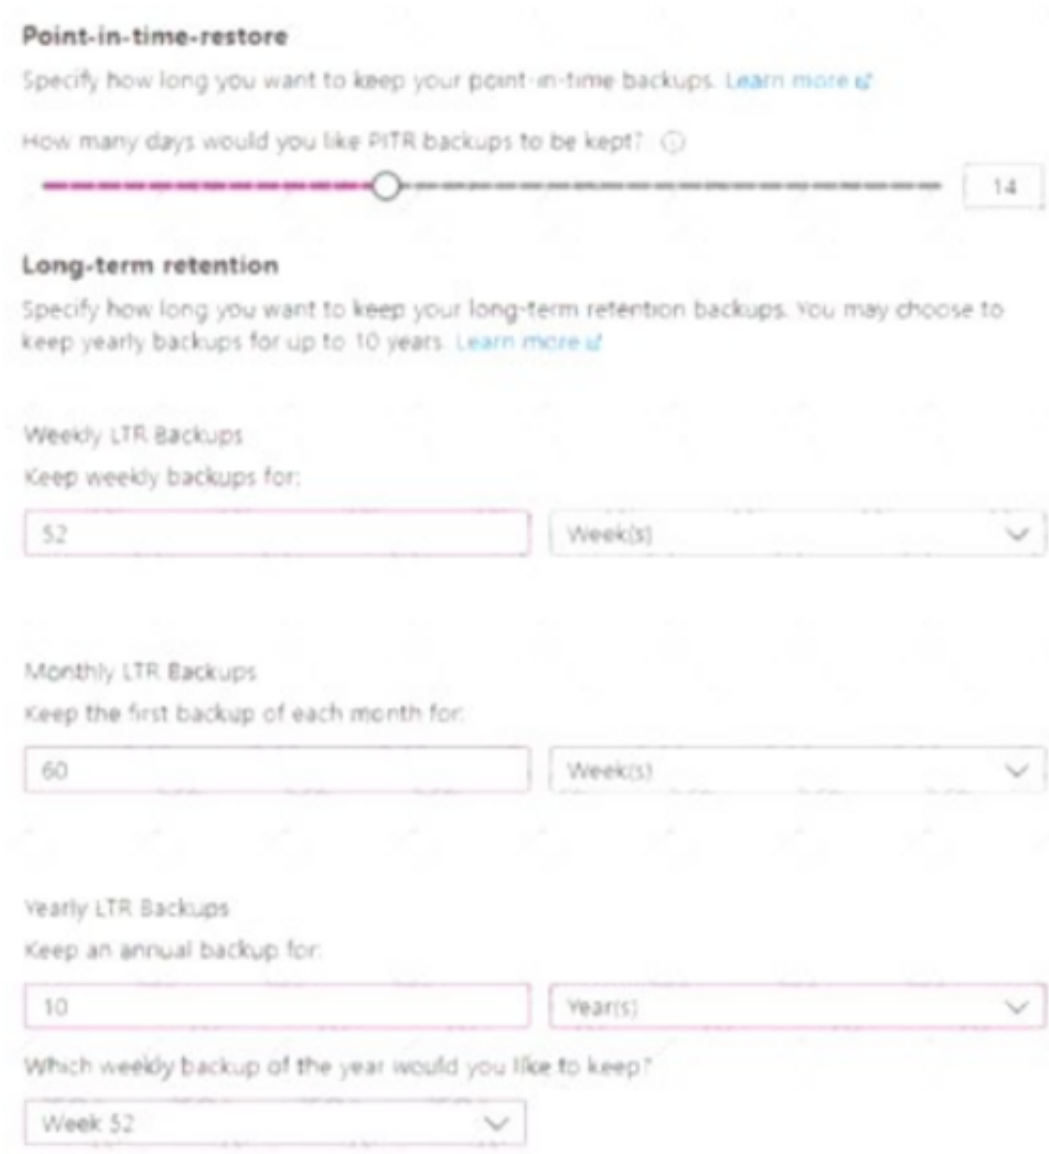

Use the drop-down menus to select the answer choice the completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

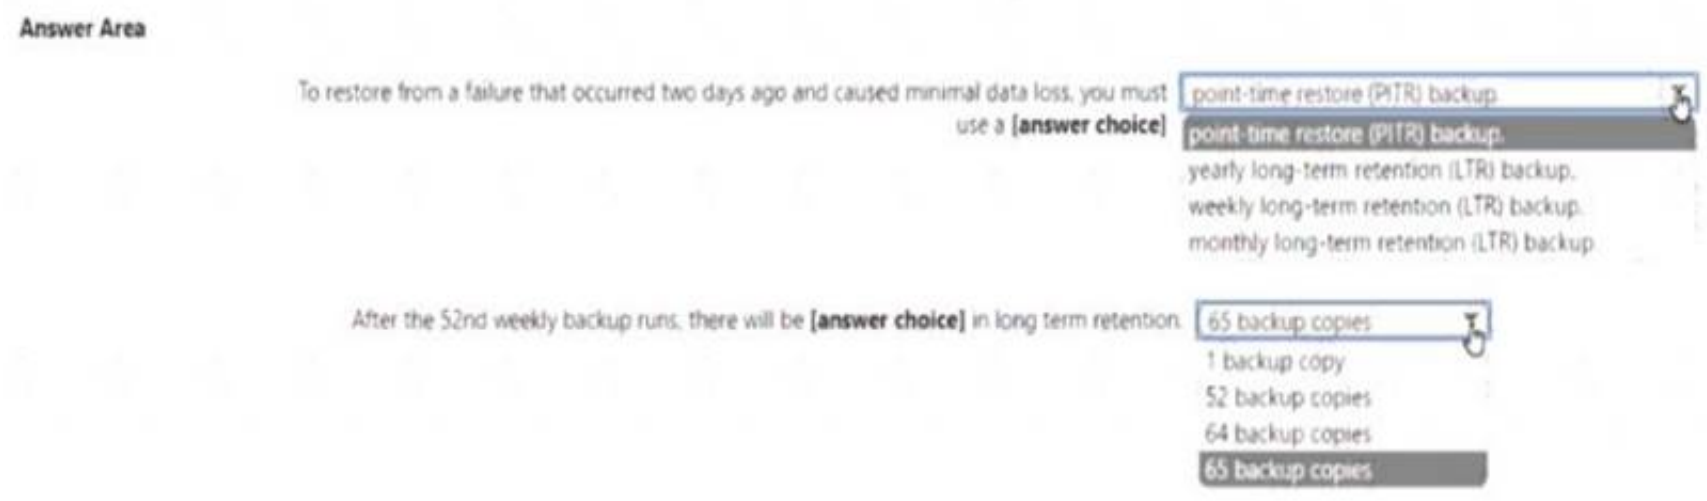

A. Mastered
B. Not Mastered

Answer: A

Explanation:

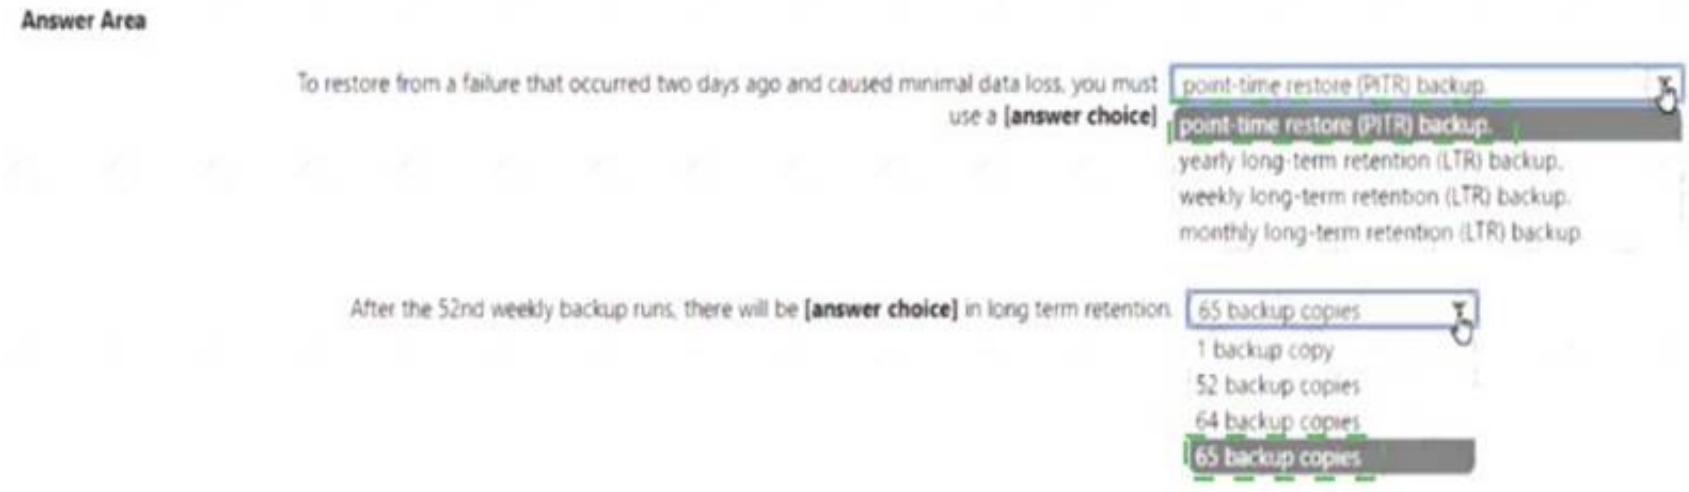

NEW QUESTION 129
- (Exam Topic 5)
You have a database on a SQL Server on Azure Virtual Machines instance.
The current state of Query Store for the database is shown in the following exhibit.

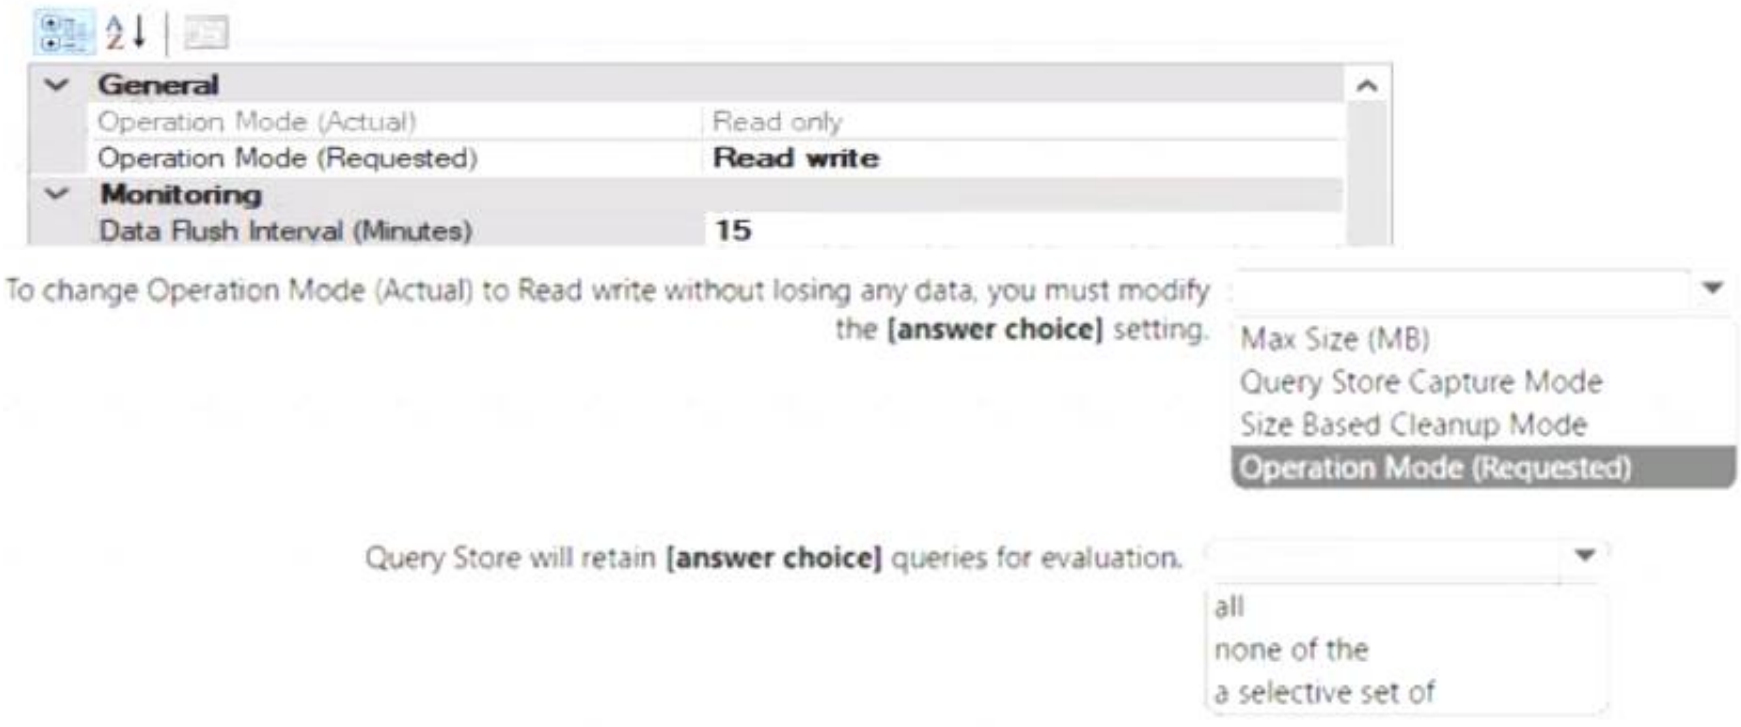

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text Description automatically generated

NEW QUESTION 134
- (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 =
'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 139
- (Exam Topic 5)
You have an Azure Synapse Analytics Apache Spark pool named Pool1.
You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.
You need to load the files into the tables. The solution must maintain the source data types. What should you do?

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

Answer: B

Explanation:
Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.
Serverless SQL pool enables you to query data in your data lake. It offers a T-SQL query surface area that accommodates semi-structured and unstructured data queries.
To support a smooth experience for in place querying of data that's located in Azure Storage files, serverless SQL pool uses the OPENROWSET function with additional capabilities.
The easiest way to see to the content of your JSON file is to provide the file URL to the OPENROWSET function, specify csv FORMAT.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-json-files https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-data-storage

NEW QUESTION 141
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to display the estimated execution plan of a query by using the query editor in the Azure portal.
What should you do first?

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

Answer: A

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

NEW QUESTION 142
- (Exam Topic 5)
You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.
ADF1 contains the following pipelines:

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

You need to configure P1 and P2 to maximize parallelism and performance.
Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

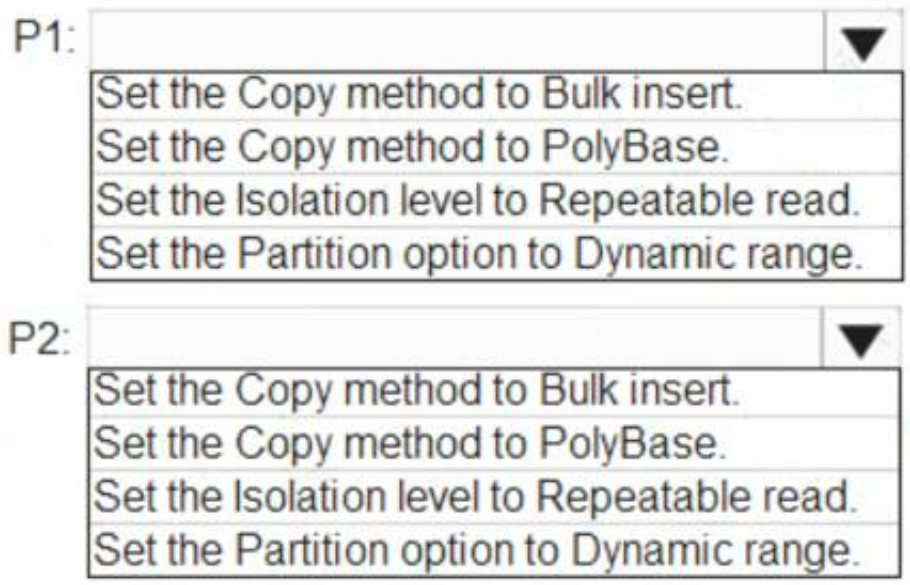

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, chat or text message Description automatically generated
P1: Set the Partition option to Dynamic Range.
The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel. P2: Set the Copy method to PolyBase
Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store. (Polybase supports Azure Blob storage and Azure Data Lake Store by default.)
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse

NEW QUESTION 145
- (Exam Topic 5)
You have an Azure SQL database.
You run the following PowerShell script.

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

Connect-AzAccount

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

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

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $serverName -
DatabaseName $dbName '
     -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention P5Y
-WeekOfYear 52
```

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

Statements	Yes	No
DB1 can be restored to a specific point in time 30 days ago.	○	○
DB1 can be restored from a weekly backup performed six months ago.	○	○
DB1 can be restored from a yearly backup performed six years ago.	○	○

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 147
- (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 149
- (Exam Topic 5)
You have an Azure subscription that contains the following resources:
• 10 Azure SQL databases
• Five Azure SQL managed instances
• Five instances of SQL Server on Azure Virtual Machines
You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort. What should you

include in the solution?

A. Log Analytics
B. Azure SQL Analytics
C. Query Performance Insight
D. SQL Insights

Answer: B

NEW QUESTION 152
- (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 157
- (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 159
- (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.

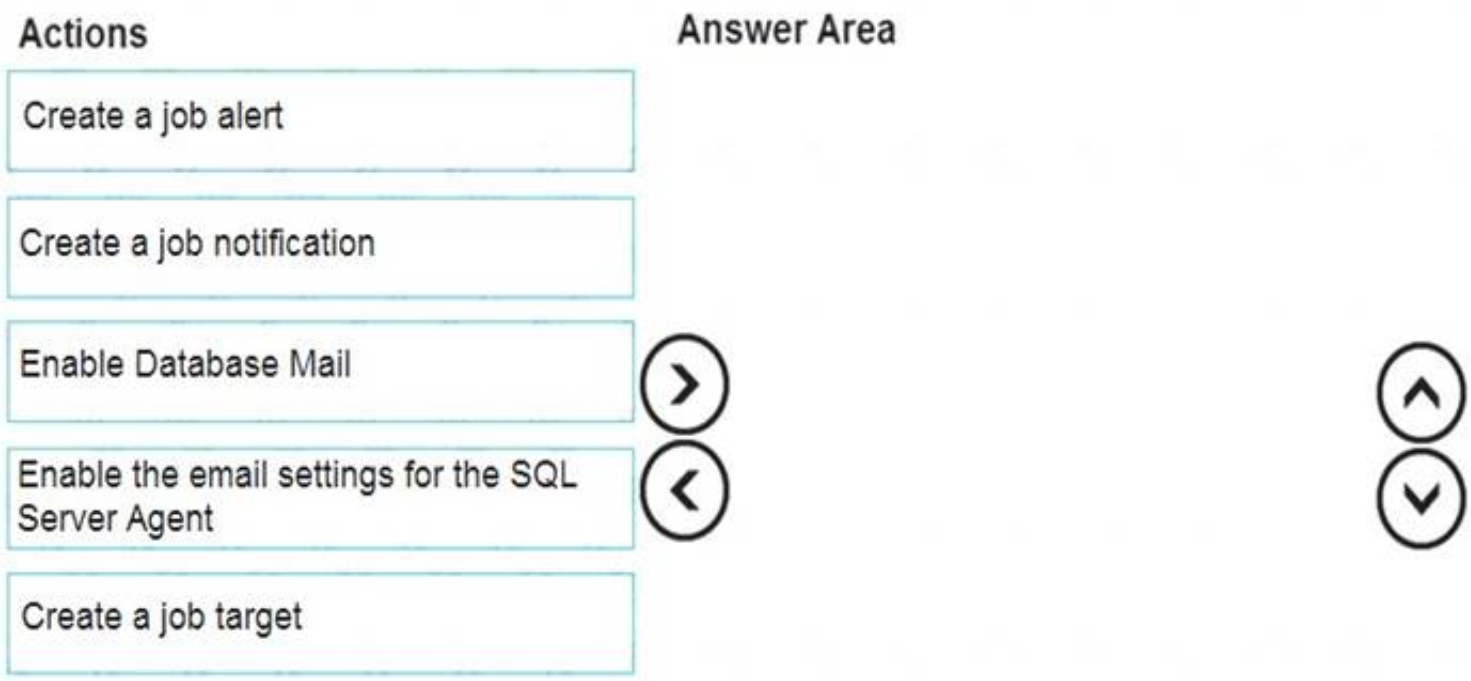

A. Mastered
B. Not Mastered

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 162
- (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 165
- (Exam Topic 5)
You have an Azure SQL managed instance.
You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries.
What should you do?

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

Answer: C

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

NEW QUESTION 166
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1. You use Query Performance Insight to monitor db1.
You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.
Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

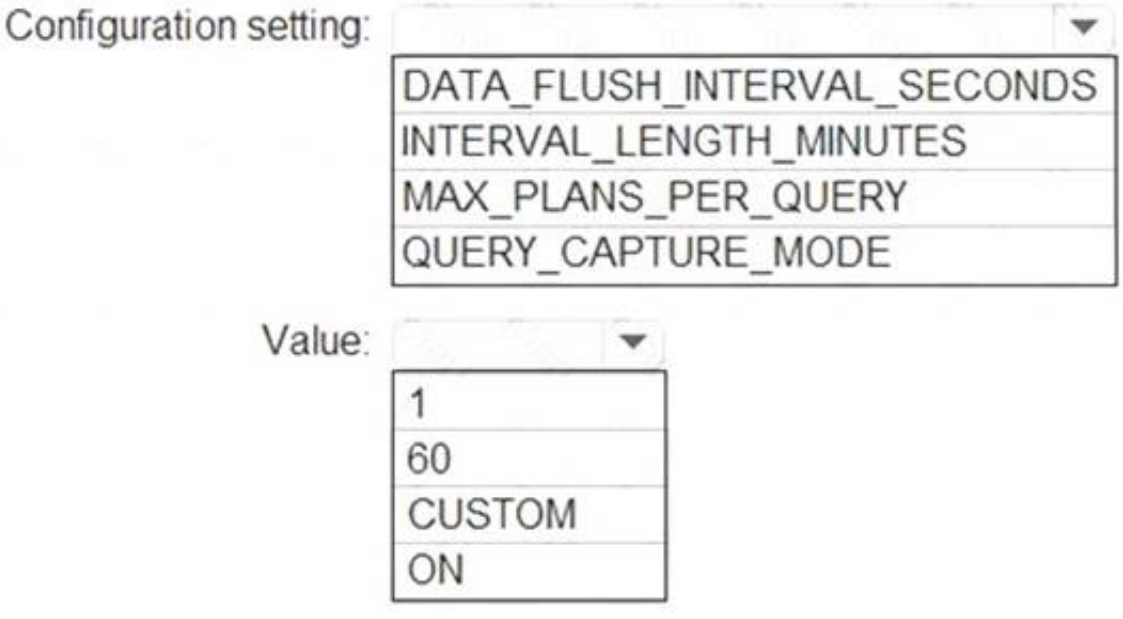

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 167
- (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 171
- (Exam Topic 5)
You have an Azure AD tenant and a logical Microsoft SQL server named SQL1 that hosts several Azure SQL databases.
You plan to assign Azure AD users permissions to the databases automatically by using Azure Automation. You need to create the required Automation accounts.
Which two accounts should you create? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Answer: AC

NEW QUESTION 174
- (Exam Topic 5)
You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1. You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.
What should you include in the recommendation?

A. Log Replay Service (LRS)
B. log shipping
C. transactional replication
D. SQL Data Sync

Answer: B

NEW QUESTION 179
- (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 181
- (Exam Topic 5)
You have an Azure SQL database named that contains a table named Table1. You run a query to bad data into Table1.
The performance Of Table1 during the load operation are shown in exhibit.

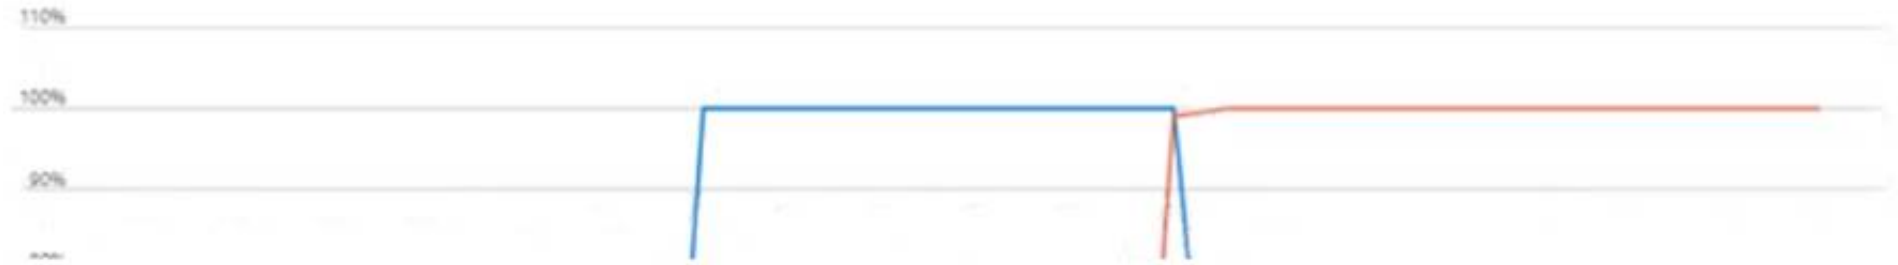

A. Mastered
B. Not Mastered

Answer: A

Explanation:

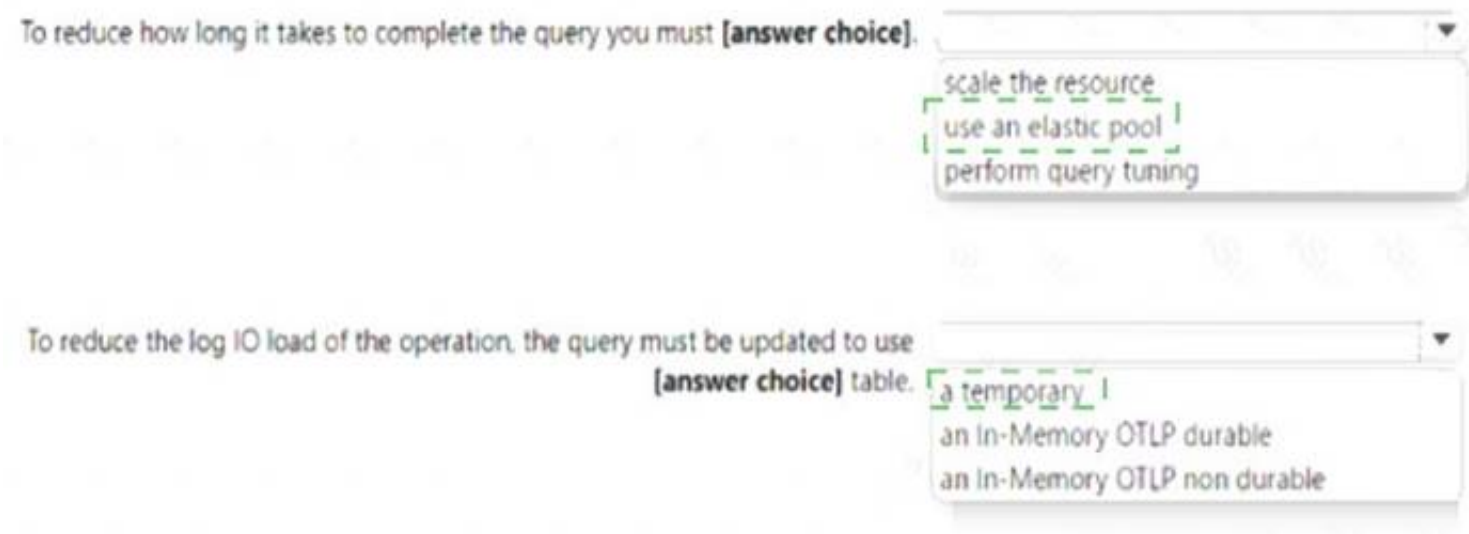

NEW QUESTION 183
- (Exam Topic 5)
You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1. You need to encrypt the data in Col1.
Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

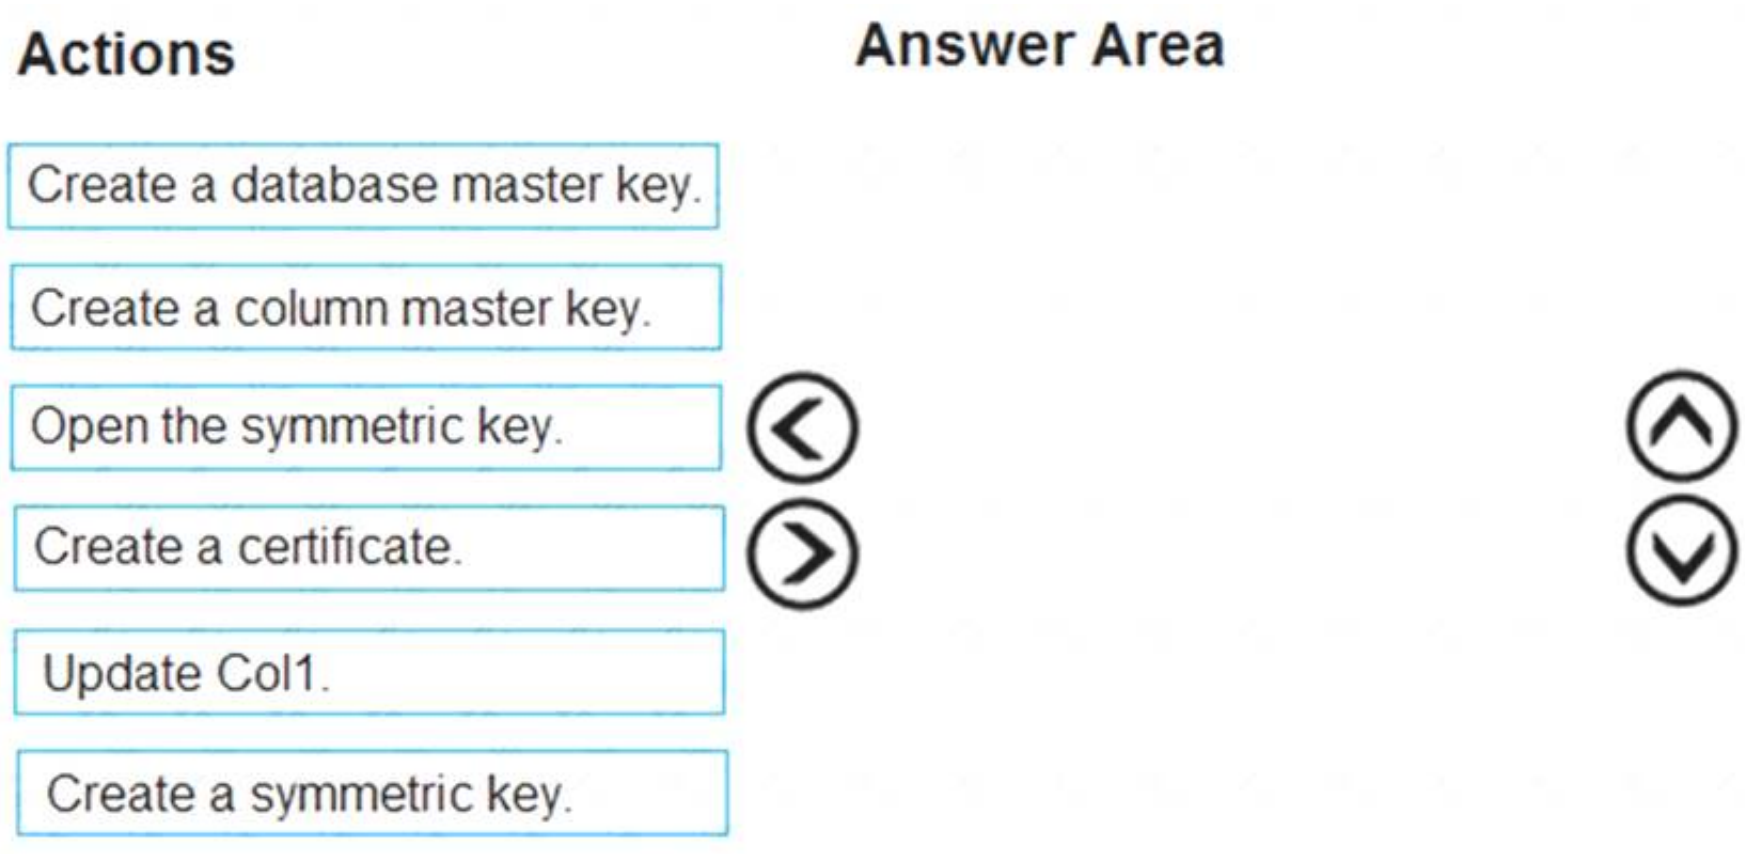

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 184
- (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 189
- (Exam Topic 5)
From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.
The data contains the following columns:

Name	Sample value
Date	15 Jan 2021
EventCategory	Videos
EventAction	Play
EventLabel	Contoso Promotional
ChannelGrouping	Social
TotalEvents	150
UniqueEvents	120
SessionsWithEvents	99

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.
To which table should you add each column? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

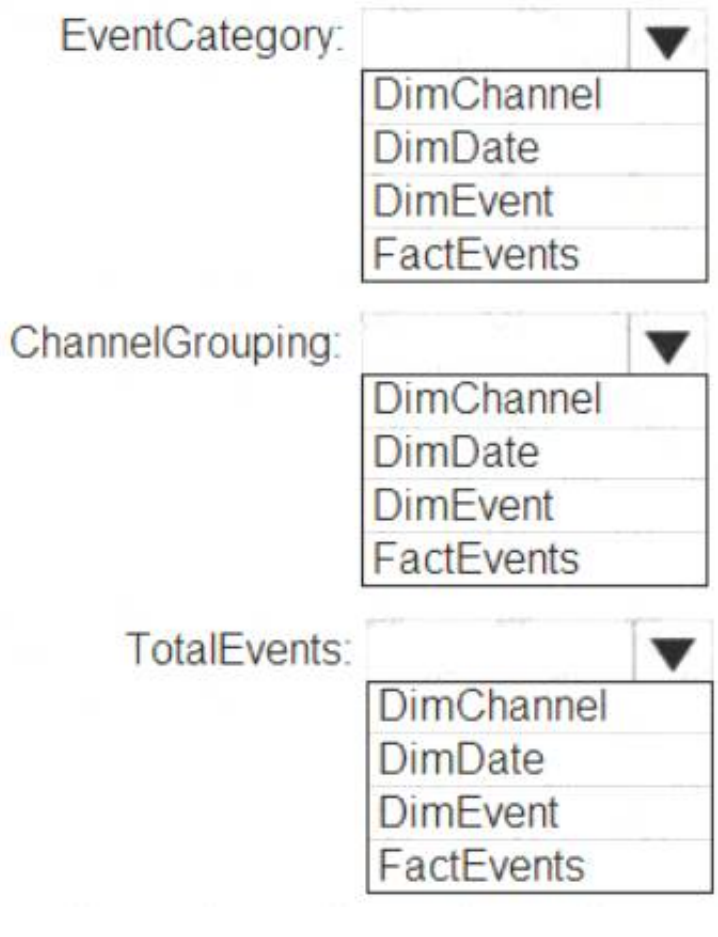

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, application, table Description automatically generated
Box 1: FactEvents
Fact tables store observations or events, and can be sales orders, stock balances, exchange rates, temperatures, etc.
Box 2: DimChannel
Dimension tables describe business entities – the things you model. Entities can include products, people, places, and concepts including time itself. The most consistent table you'll find in a star schema is a date dimension table. A dimension table contains a key column (or columns) that acts as a unique identifier, and descriptive columns.
Box 3: DimEvent Reference:
https://docs.microsoft.com/en-us/power-bi/guidance/star-schema

NEW QUESTION 192
- (Exam Topic 5)
You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.
Data to be loaded is identified by a column named LastUpdatedDate in the source table. You plan to execute the pipeline every four hours.
You need to ensure that the pipeline execution meets the following requirements:
Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.
Which type of trigger should you use?

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

Answer: A

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

NEW QUESTION 195
- (Exam Topic 2)
You are evaluating the role assignments.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE:Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 198
- (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 202
- (Exam Topic 1)
You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

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

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: DTU
Scenario:

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

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

NEW QUESTION 204
- (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.

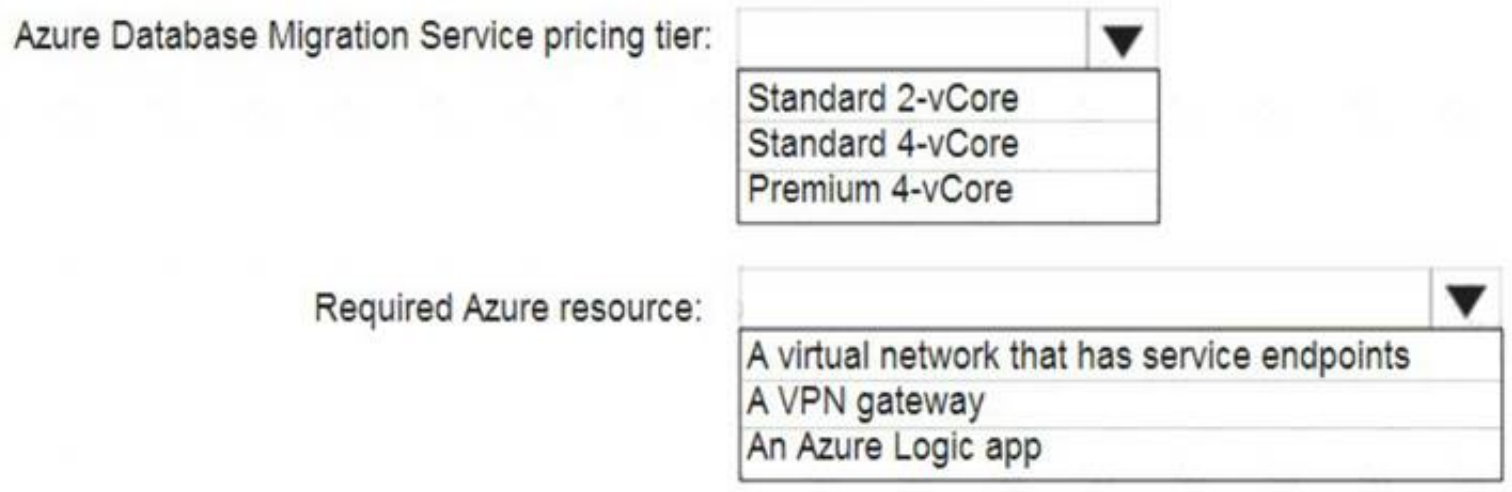

A. Mastered
B. Not Mastered

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 209
- (Exam Topic 1)
You create all of the tables and views for ResearchDB1.
You need to implement security for ResearchDB1. The solution must meet the security and compliance requirements.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

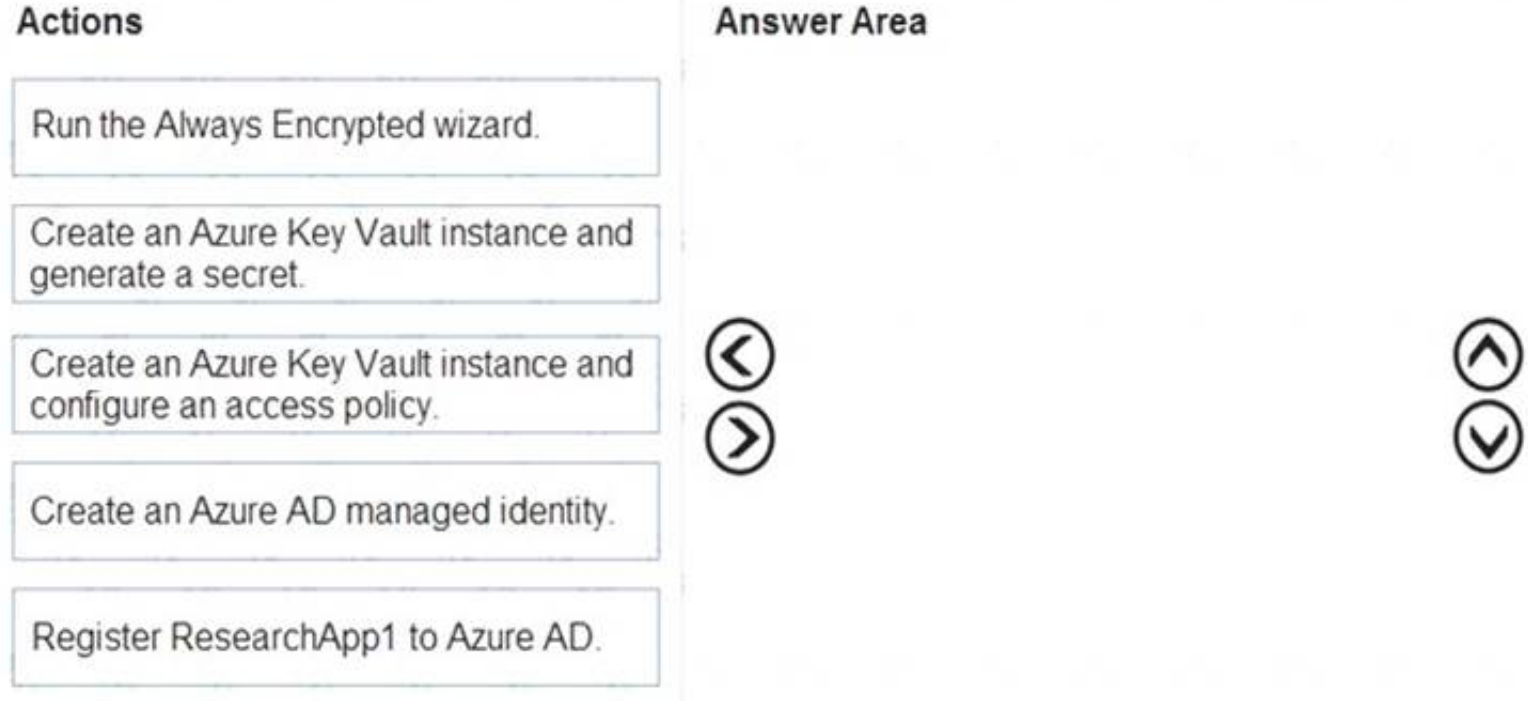

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 214
- (Exam Topic 1)
You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance requirements.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

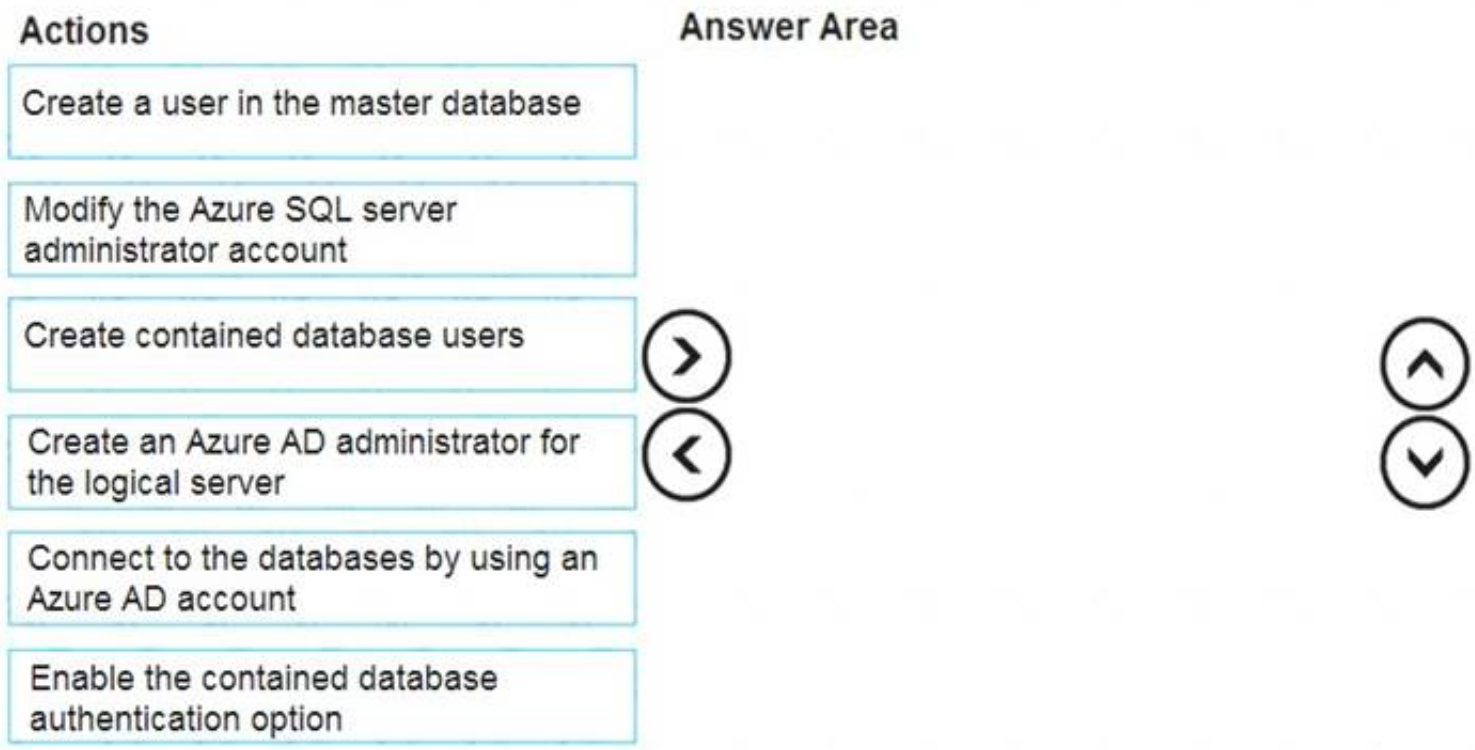

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Authenticate database users by using Active Directory credentials.
The configuration steps include the following procedures to configure and use Azure Active Directory authentication.
- Create and populate Azure AD.
- Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
- Create an Azure Active Directory administrator. (Step 1)
- Configure your client computers.
- Create contained database users in your database mapped to Azure AD identities. (Step 2)
- Connect to your database by using Azure AD identities. (Step 3)

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview

NEW QUESTION 217
- (Exam Topic 1)
You need to identify the cause of the performance issues on SalesSQLDb1.
Which two dynamic management views should you use? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: AE

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

sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-lock

NEW QUESTION 221
- (Exam Topic 1)
You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.
What should you include in the recommendation? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

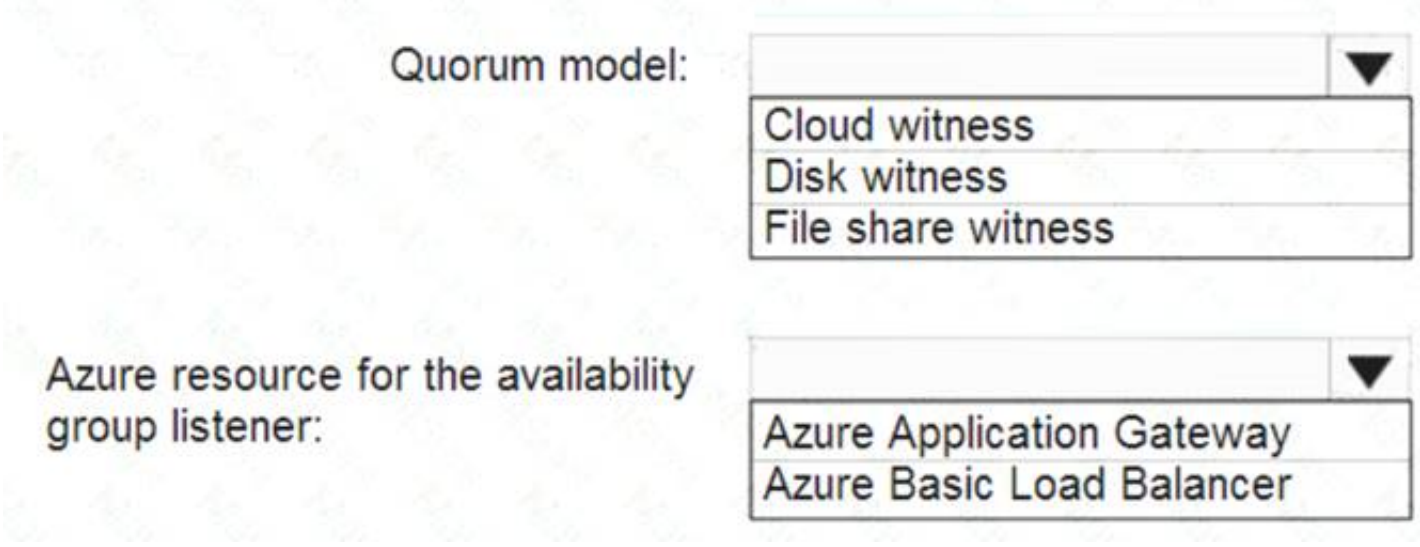

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Business Requirements
Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.
Box 1: Cloud witness
If you have a Failover Cluster deployment, where all nodes can reach the internet (by extension of Azure), it is recommended that you configure a Cloud Witness as your quorum witness resource.
Box 2: Azure Basic Load Balancer
Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.
Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a publicVirtual IP (VIP) that is accessible over the internet. An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.
Reference:
https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

NEW QUESTION 225
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DP-300 Practice Exam Features:

* DP-300 Questions and Answers Updated Frequently

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

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

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