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 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.

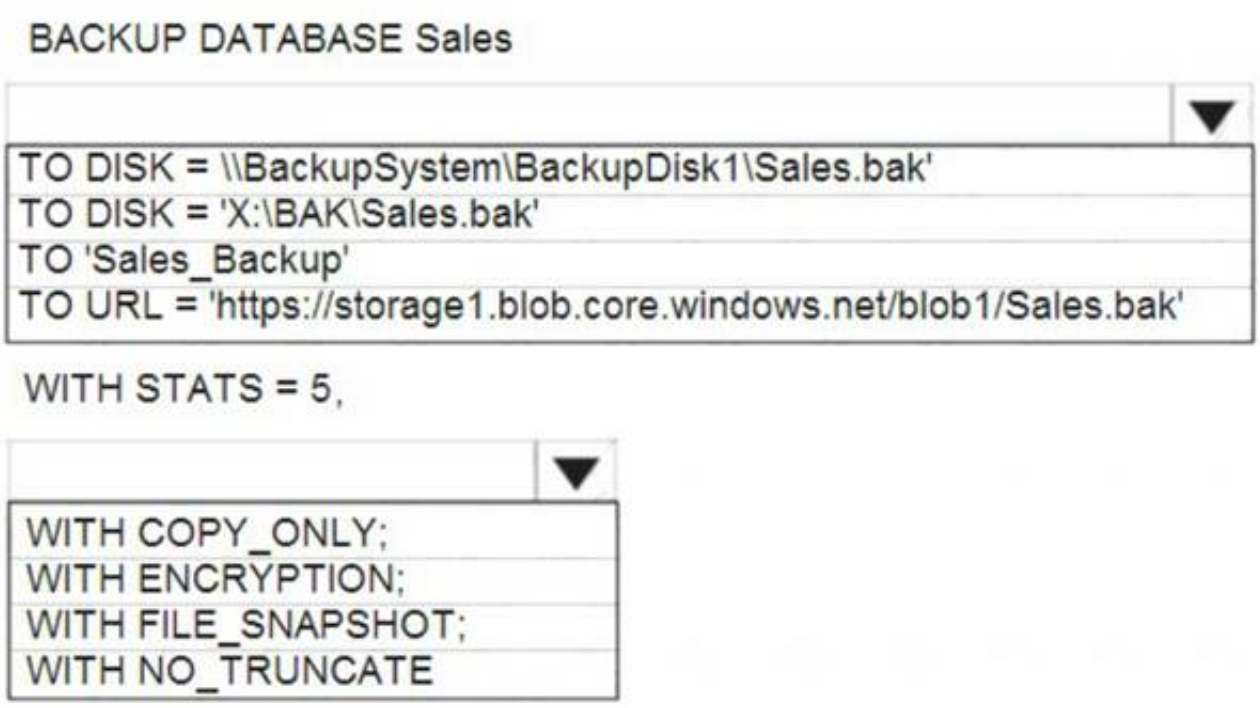

A. Mastered
B. Not Mastered

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 SQL Server on an Azure virtual machine that contains a database named DB1. The database reports a CHECKSUM error.
You need to recover the database.
How should you complete the statements? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

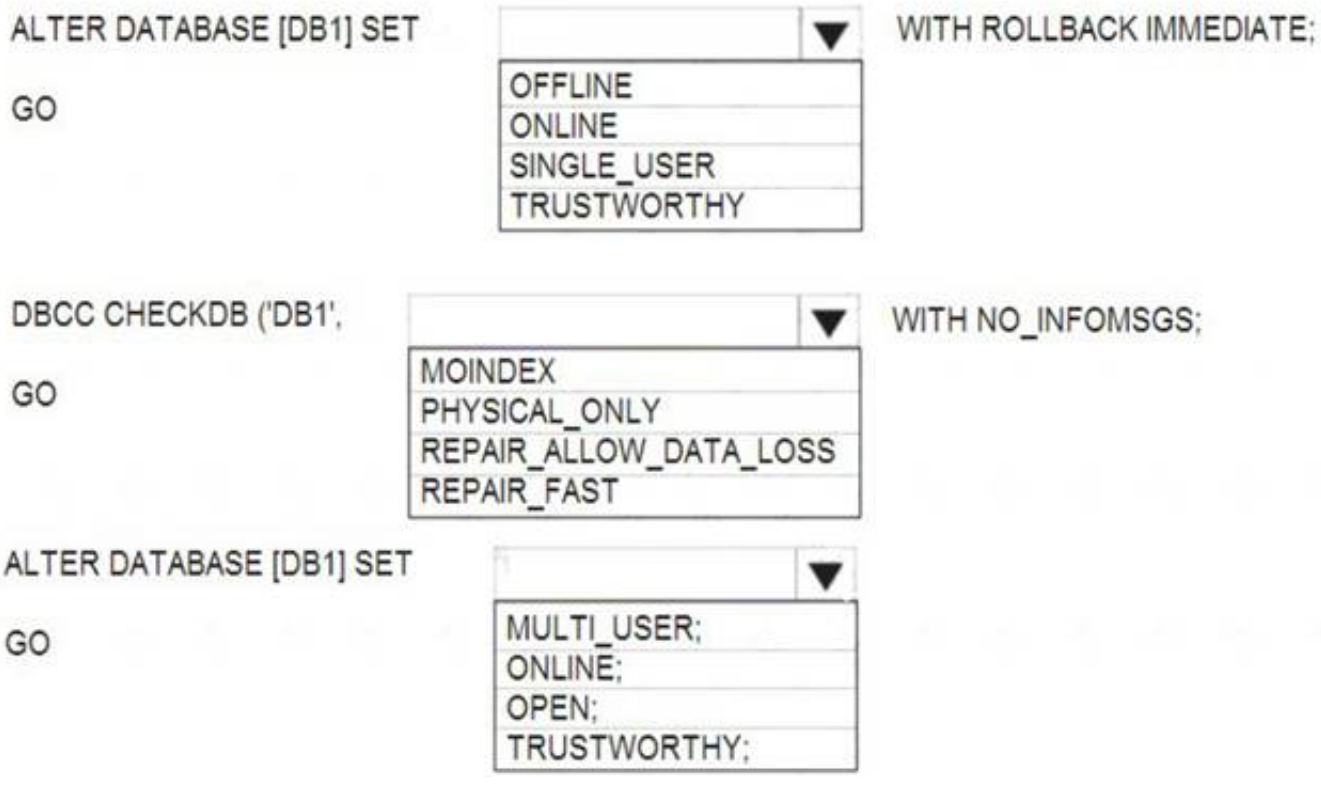

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SINGLE_USER
The specified database must be in single-user mode to use one of the following repair options. Box 2: REPAIR_ALLOW_DATA_LOSS
REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.
Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql

NEW QUESTION 4
- (Exam Topic 5)
You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables.
Which distribution type should you recommend to minimize data movement?

A. HASH
B. REPLICATE
C. ROUND_ROBIN

Answer: B

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

NEW QUESTION 5
- (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 6
- (Exam Topic 5)
You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.
Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted. You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the
amount of time required to delete old data.
How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.
You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

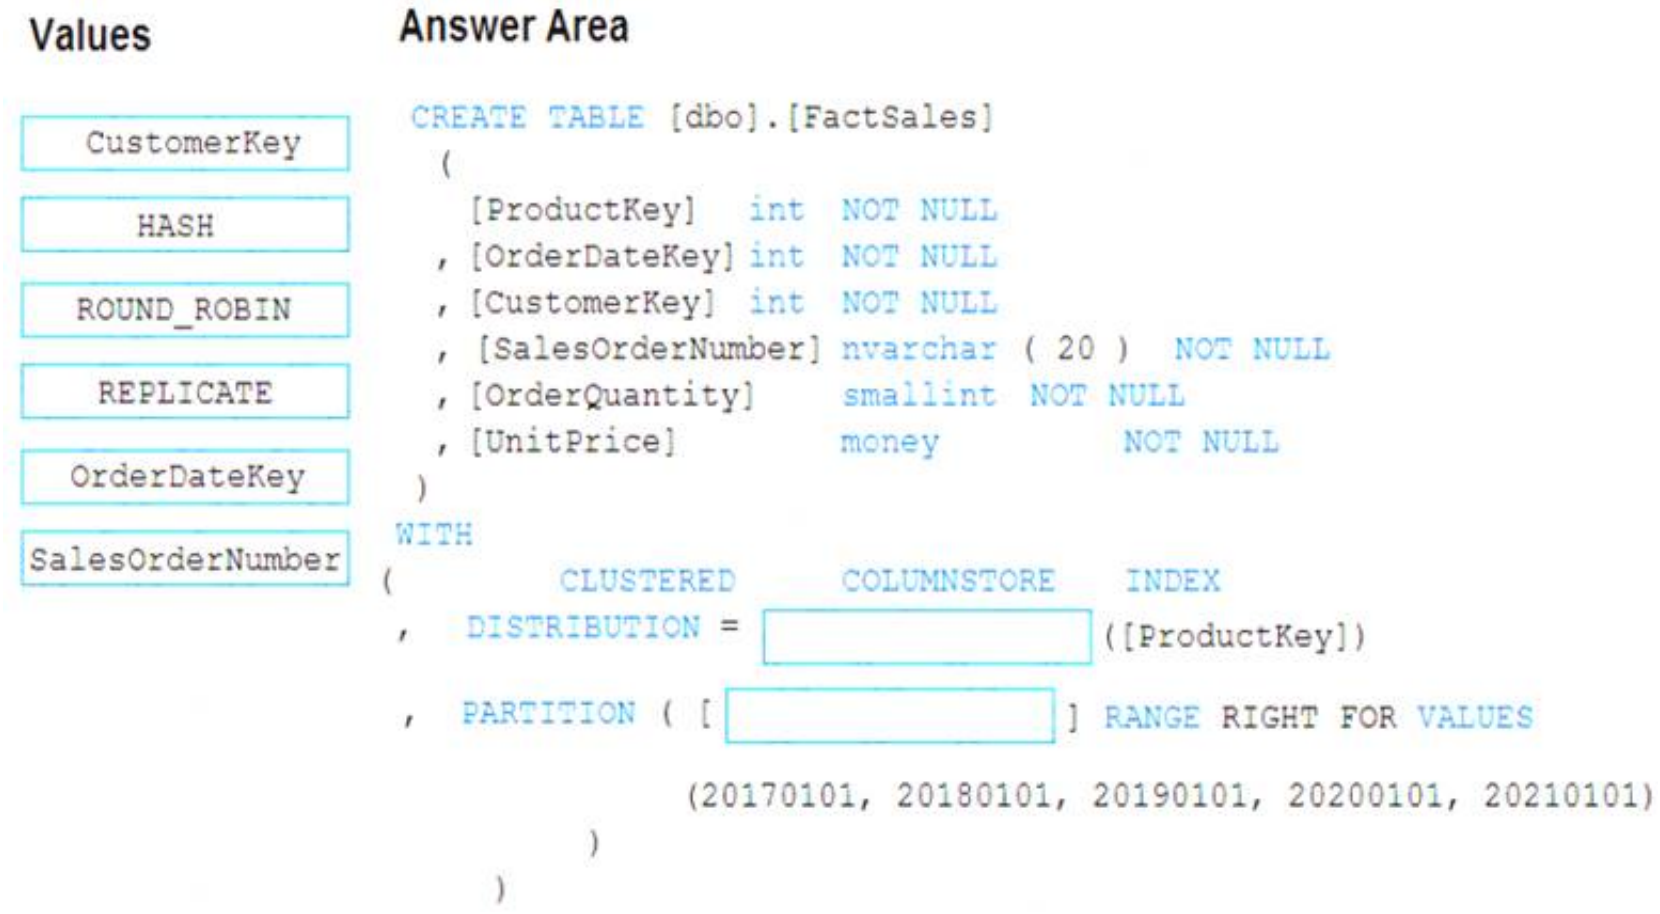

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 7
- (Exam Topic 5)
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

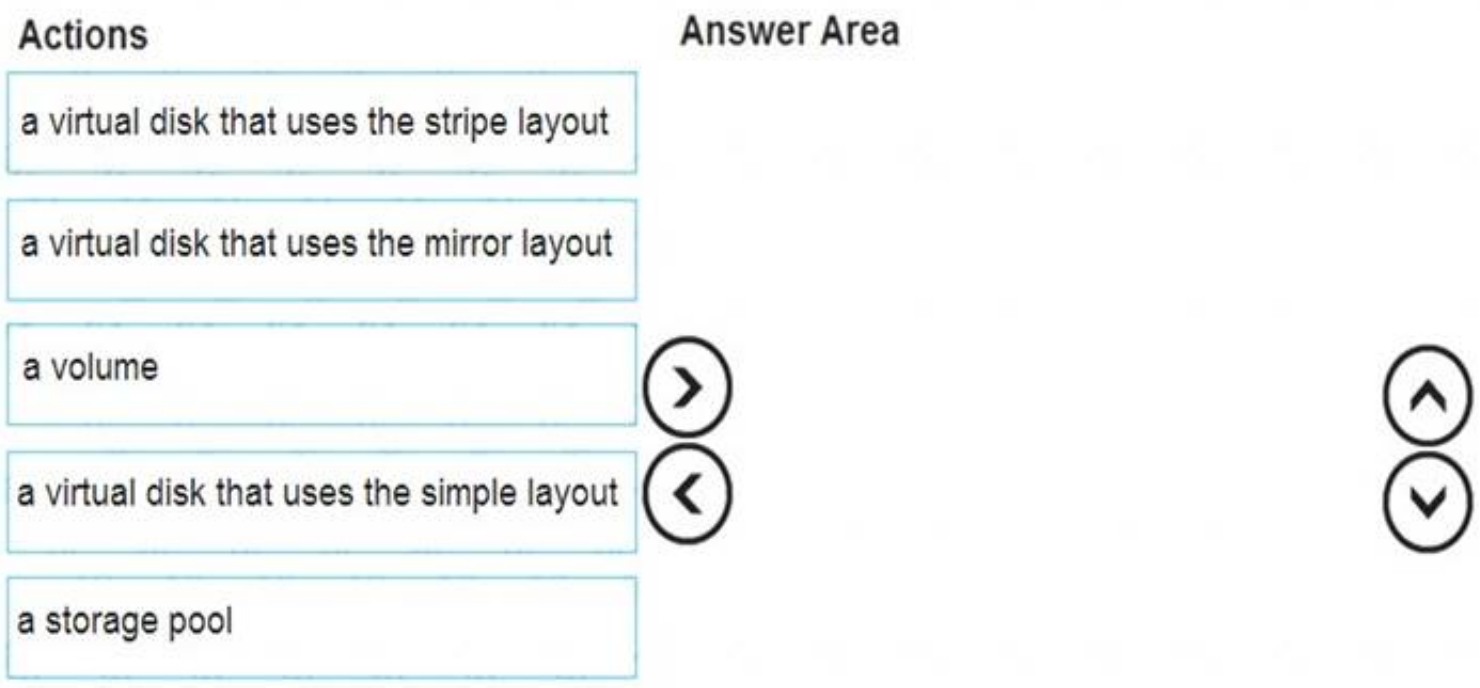

A. Mastered
B. Not Mastered

Answer: A

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

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

NEW QUESTION 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 SQL Server on Azure virtual machines in an availability group. You have a database named DB1 that is NOT in the availability group.
You create a full database backup of DB1. You need to add DB1 to the availability group.
Which restore option should you use on the secondary replica?

A. Restore with Recovery

B. Restore with Norecovery
C. Restore with Standby

Answer: B

Explanation:
Prepare a secondary database for an Always On availability group requires two steps:
* 1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY
* 2. Join the restored database to the availability group. Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-asecondary database-for-an-availability-group-sql-server

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 11
- (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.

Answer Area

Query Store will retain **[answer choice]** queries for evaluation.

To change Operation Mode (Actual) to Read write without losing any data, you must modify the **[answer choice]** setting.

To change Operation Mode (Actual) to Read write without losing any data, you must modify the **[answer choice]** setting.

Max Size (MB)
Query Store Capture Mode
Size Based Cleanup Mode
Operation Mode (Requested)

Query Store will retain **[answer choice]** queries for evaluation.

all
none of the
a selective set of

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text Description automatically generated

NEW QUESTION 16
- (Exam Topic 5)
You are developing an application that uses Azure Data Lake Storage Gen 2.
You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Answer: C

Explanation:
A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:
What resources the client may access.
What permissions they have to those resources. How long the SAS is valid.
Note: Data Lake Storage Gen2 supports the following authorization mechanisms:

- Shared Key authorization
- Shared access signature (SAS) authorization
- Role-based access control (Azure RBAC)
- Shared Key authorization
- Shared access signature (SAS) authorization
- Role-based access control (Azure RBAC)
- Access control lists (ACL)

Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 20
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.
You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.
Solution: You restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
The REPLACE option overrides several important safety checks that restore normally performs. The overridden checks are as follows:

- Restoring over an existing database with a backup taken of another database.

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

NEW QUESTION 21
- (Exam Topic 5)
You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.
Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure. You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.
What should you include in the recommendation?

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

Answer: D

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

NEW QUESTION 24
- (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.

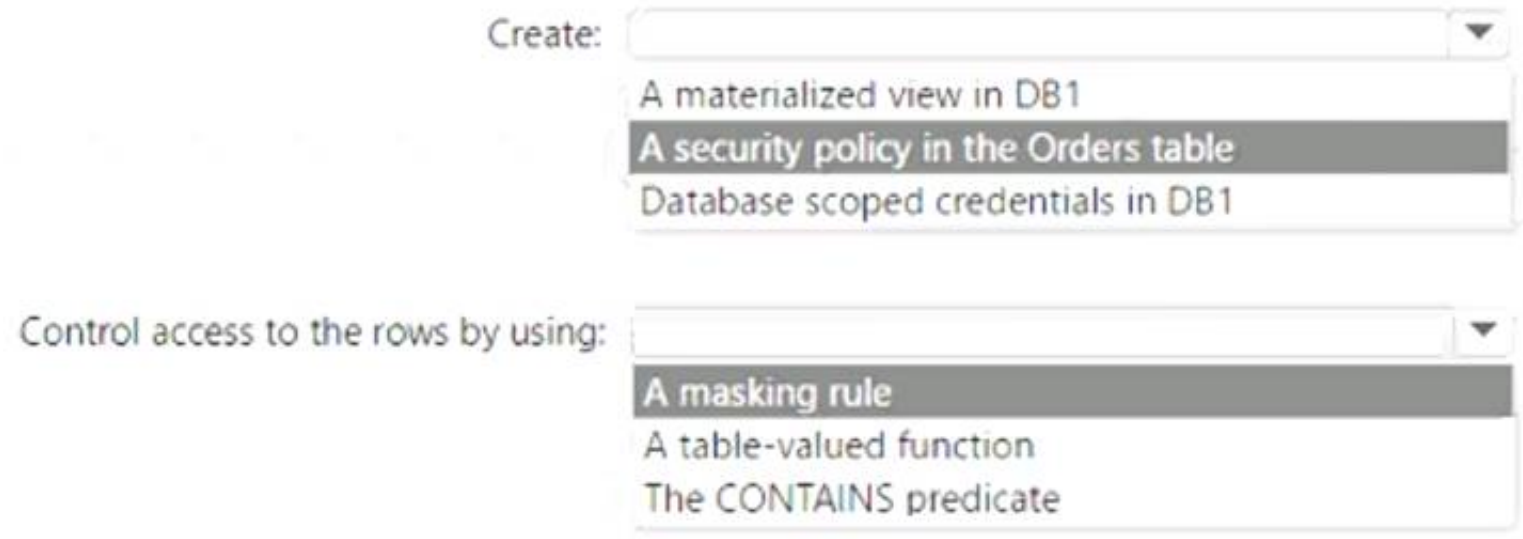

A. Mastered
B. Not Mastered

Answer: A

Explanation:

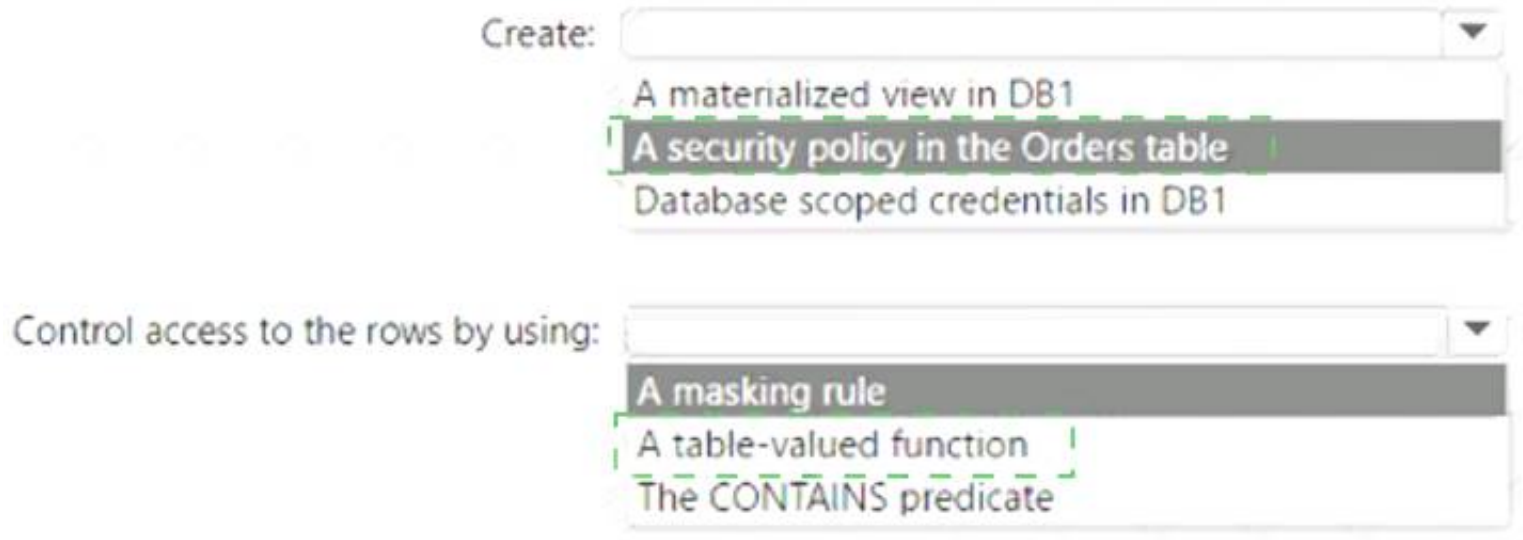

NEW QUESTION 27
- (Exam Topic 5)
You have an Azure SQL database named sqldb1.
You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

A. Double the value of Data Flush interval
B. Decrease by half the value of Data Flush Interval
C. Double the value of Statistics Collection Interval
D. Decrease by half the value of Statistics Collection interval

Answer: B

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

NEW QUESTION 29
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1.
You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:

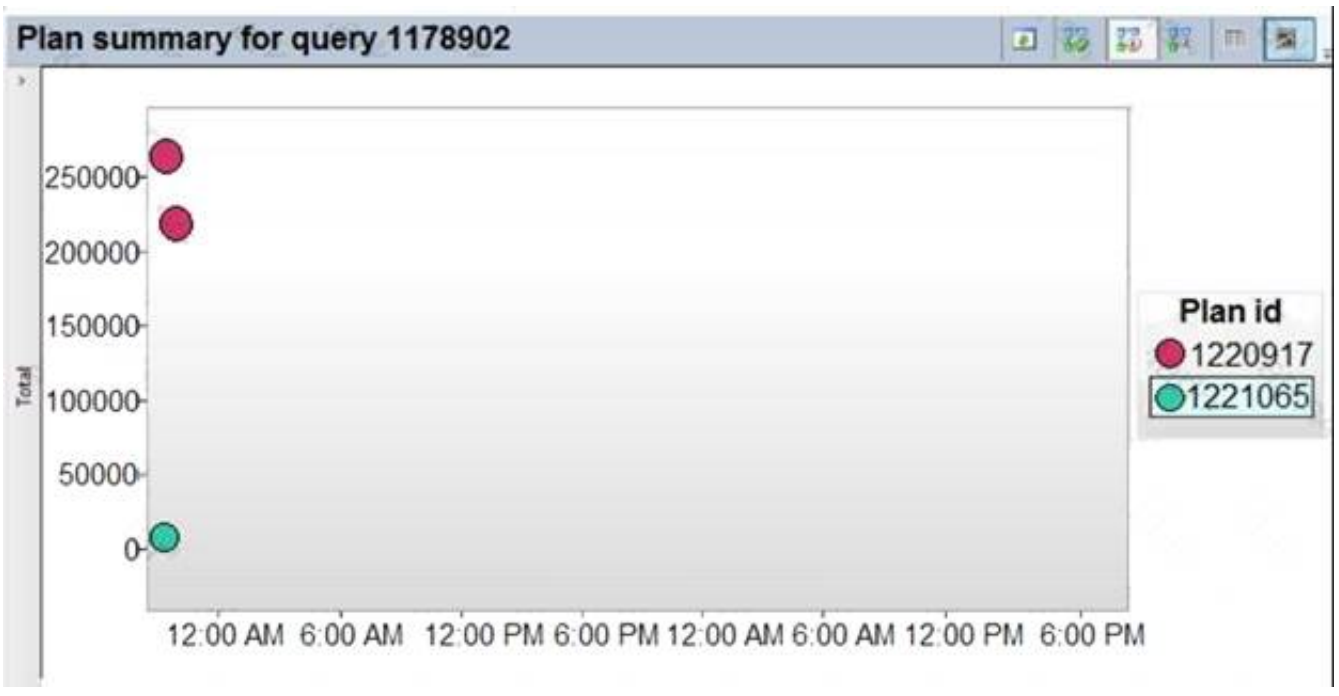

What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

A. Force the query execution plan for plan 1221065.
B. Run the DBCC FREEPROCCACHE command.
C. Force the query execution plan for plan 1220917.
D. Disable parameter sniffing.

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios

NEW QUESTION 33
- (Exam Topic 5)
A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate    DATETIME NULL,
    StreetAddress1  VARCHAR(500) NOT NULL,
    StreetAddress2  VARCHAR(500) NOT NULL,
    StreetAddress3  VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Change Salary to the money data type.
B. Change PhoneNumber to the float data type.
C. Change LastHireDate to the datetime2(7) data type.
D. Change PhoneNumber to the bigint data type.
E. Change LastHireDate to the date data type.

Answer: AE

NEW QUESTION 36
- (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 38
- (Exam Topic 5)
You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the Exhibit tab.)

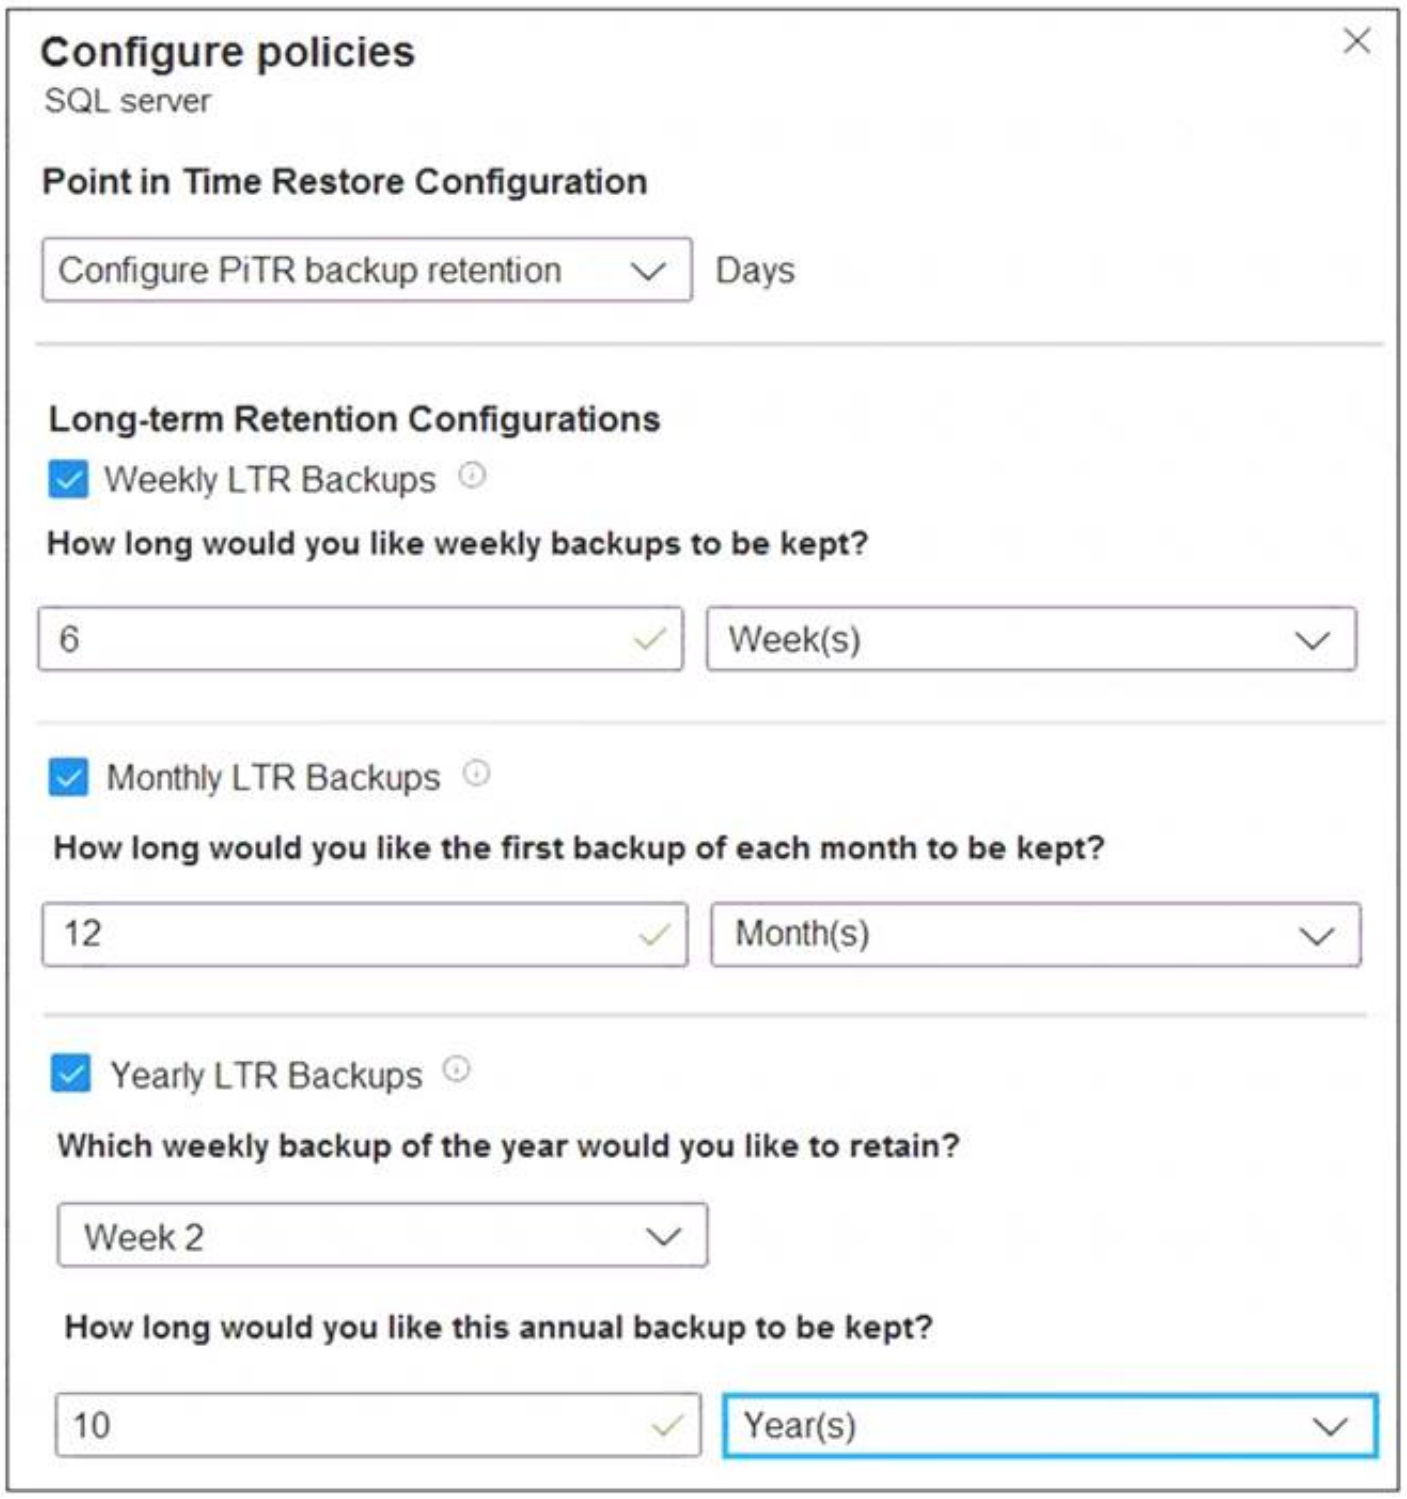

The first weekly backup occurred on January 4, 2020. The dates for the first 10 weekly backups are:
- January 4, 2020
- January 11, 2020
- January 18, 2020
- January 25, 2020
- February 1, 2020
- February 8, 2020
- February 15, 2020
- February 22, 2020
- February 29, 2020
- March 7, 2020

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.

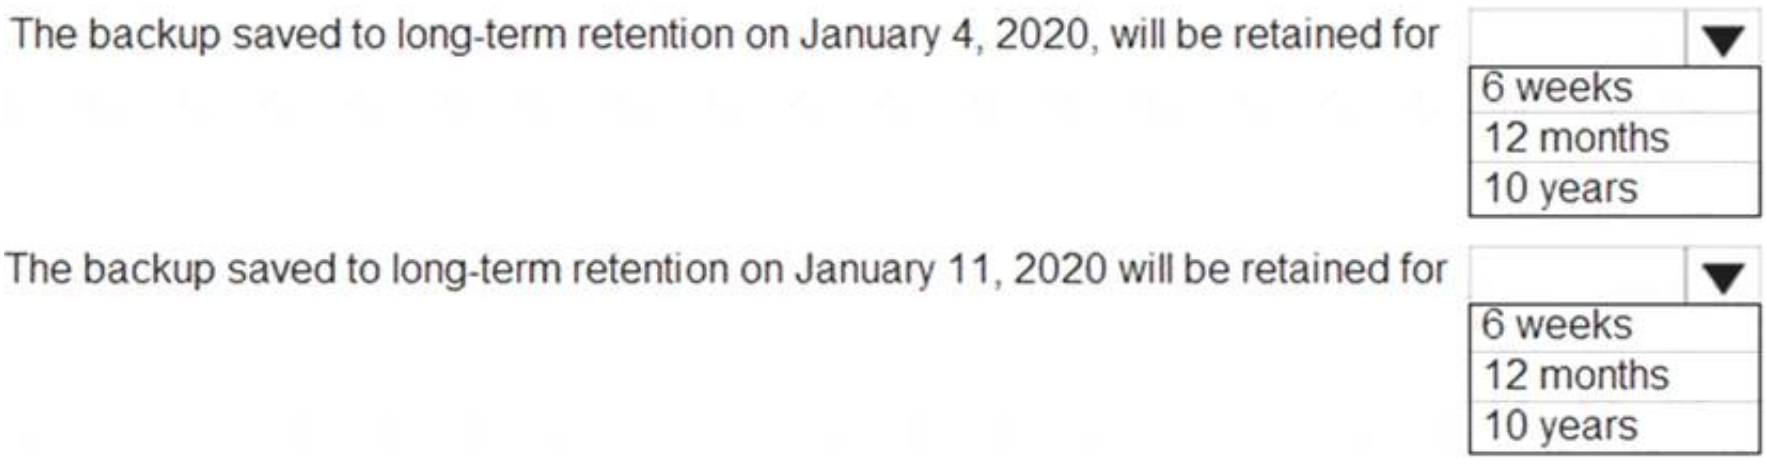

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 43
- (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 47
- (Exam Topic 5)
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.
You scale up the virtual machine to 8 vCPUSs and 64 GB of memory. You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?

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

Answer: C

Explanation:
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

NEW QUESTION 49
- (Exam Topic 5)
You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database. A recent failure causes the master database to be lost.
You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.
Which four actions should you perform in sequence to resolve the issue? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct.

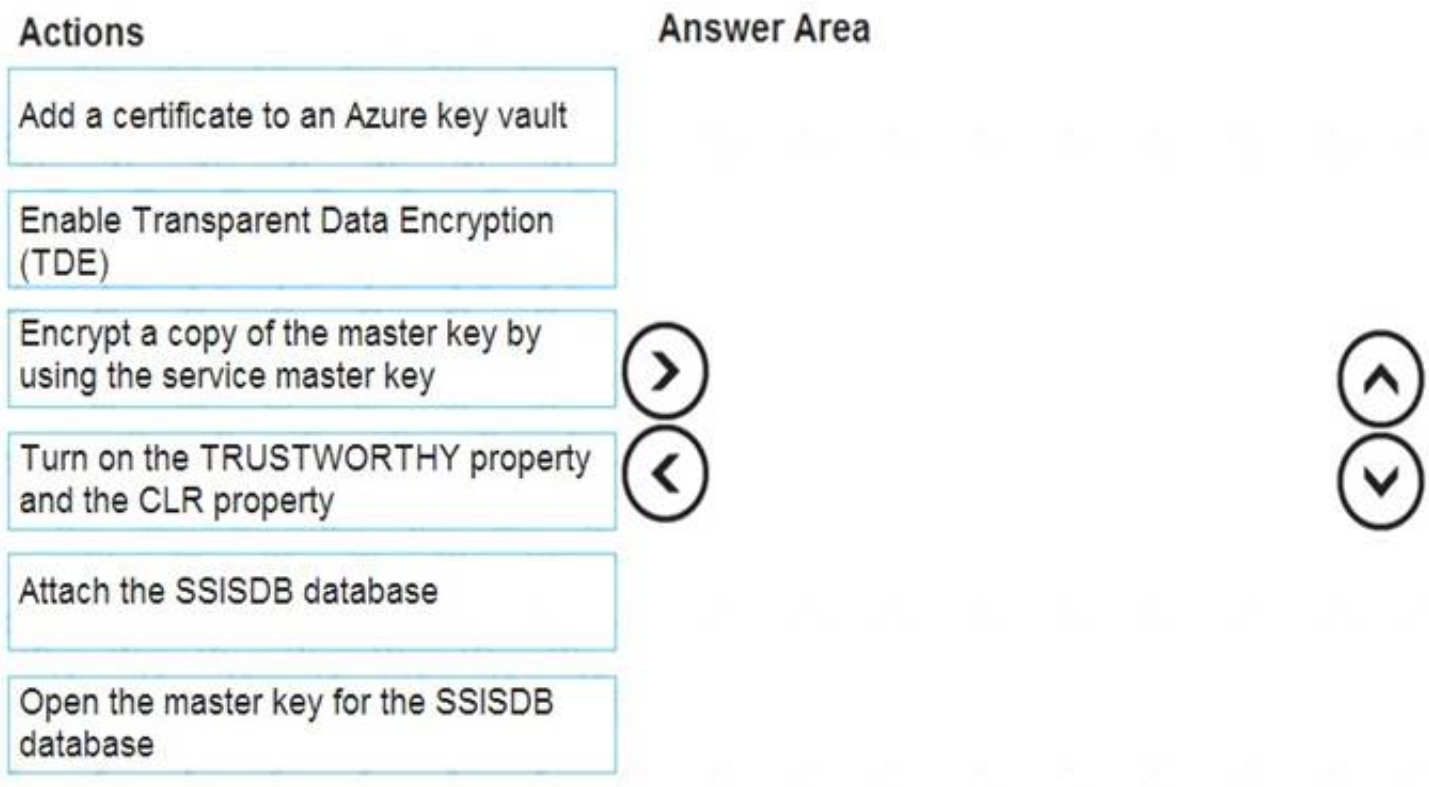

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 54
- (Exam Topic 5)
You have an Azure SQL database that contains a table named factSales. FactSales contains the columns shown in the following table.

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

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

A. page compression
B. row compression
C. columnstore compression
D. columnstore archival compression

Answer: D

Explanation:
Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.
Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression

NEW QUESTION 57
- (Exam Topic 5)
You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.
Which resource provider should you enable?

A. Microsoft.EventHub
B. Microsoft.EventGrid
C. Microsoft.Sql

D. Microsoft.Automation

Answer: B

Explanation:
Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger

NEW QUESTION 60
- (Exam Topic 5)
You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

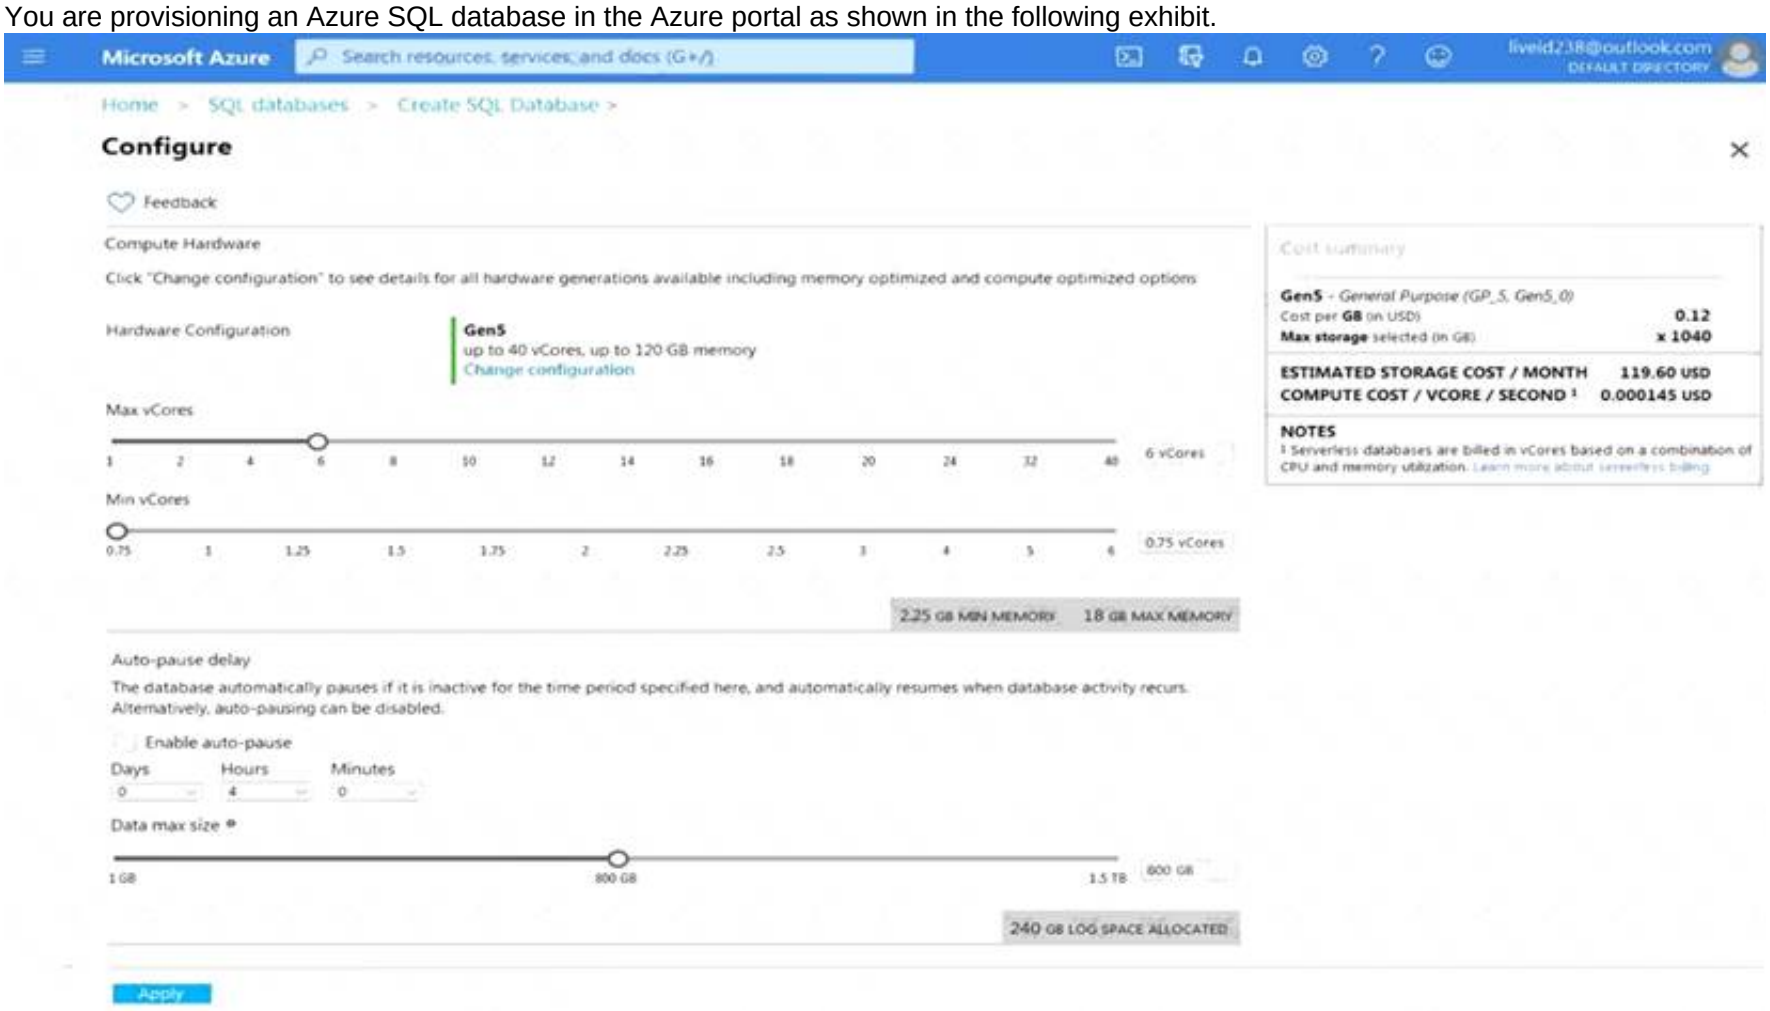

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

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

- no extra time
- up to 10 minutes
- up to one minute

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

- intermittent and unpredictable
- regular and high
- steady and low

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 62
- (Exam Topic 5)
You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.
You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.
Which resource type should you use to create the alert?

A. Resource Groups
B. SQL Servers
C. SQL Databases
D. SQL Virtual Machines

Answer: C

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

NEW QUESTION 66
- (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 68
- (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 69
- (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.

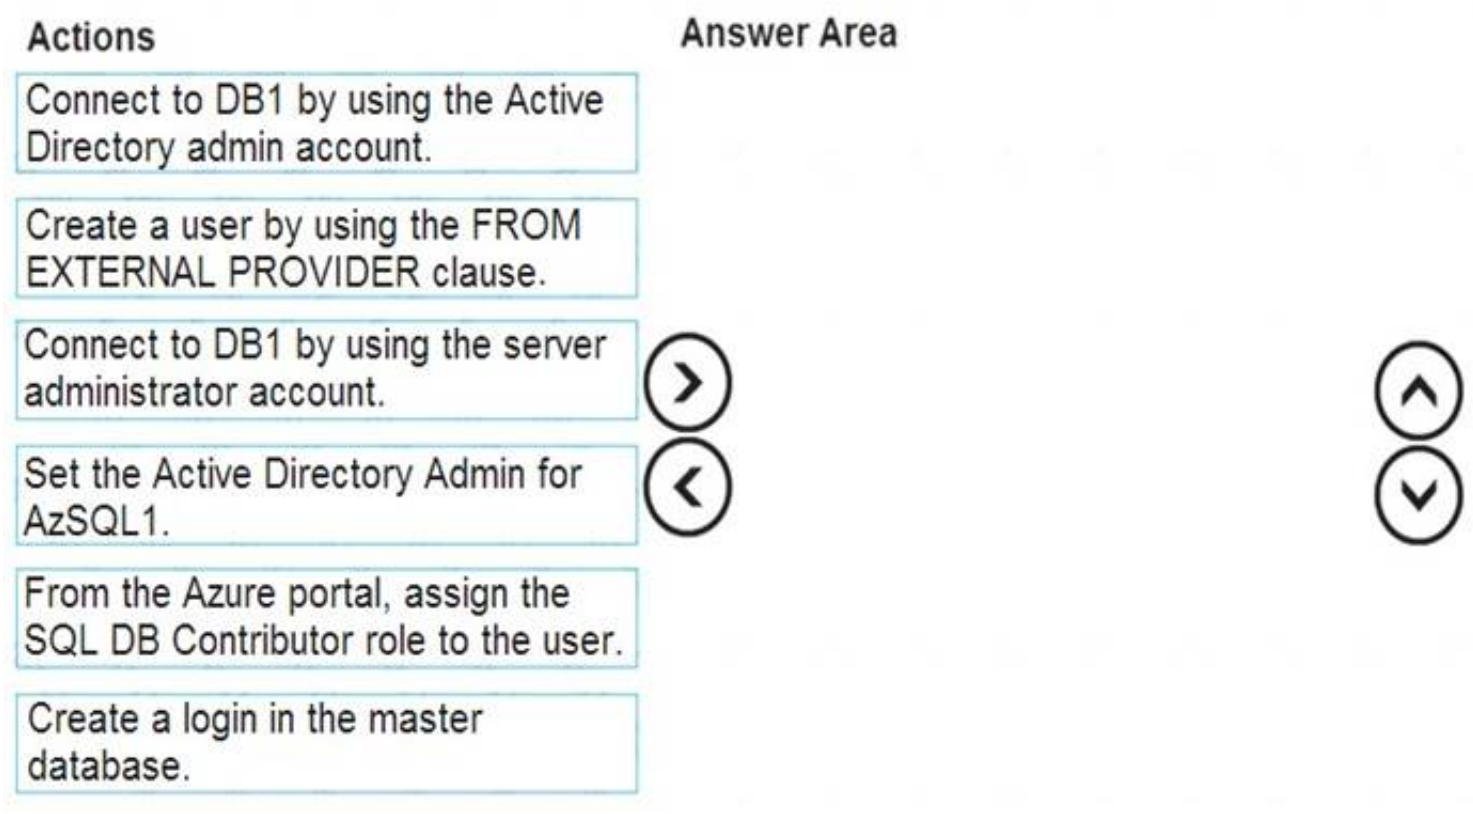

A. Mastered
B. Not Mastered

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 72
- (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 75
- (Exam Topic 5)
You have an Azure SQL database.
You run the following PowerShell script.

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

Connect-AzAccount

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

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

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

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

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

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 78
- (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 80
- (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 85
- (Exam Topic 5)
You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.
You plan to have a fact table for website visits. The table will be approximately 5 GB.
You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

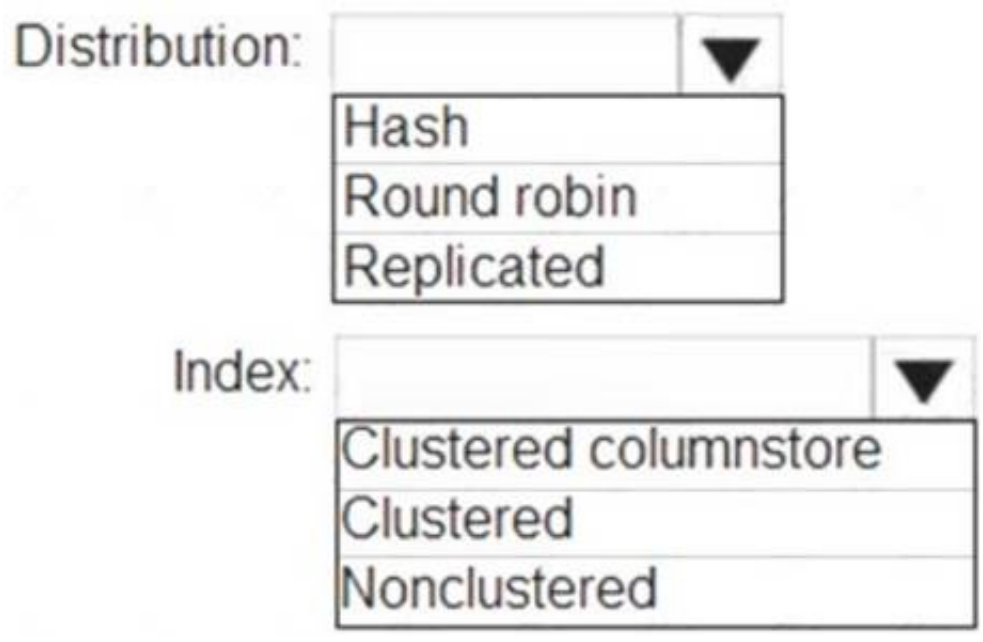

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, table, chat or text message Description automatically generated
Box 1: Hash
Consider using a hash-distributed table when:
The table size on disk is more than 2 GB.

The table has frequent insert, update, and delete operations. Box 2: Clustered columnstore
Clustered columnstore tables offer both the highest level of data compression and the best overall query performance.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-index

NEW QUESTION 88
- (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 89
- (Exam Topic 5)
You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

A. automated
B. interactive
C. High Concurrency

Answer: A

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

NEW QUESTION 94
- (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.

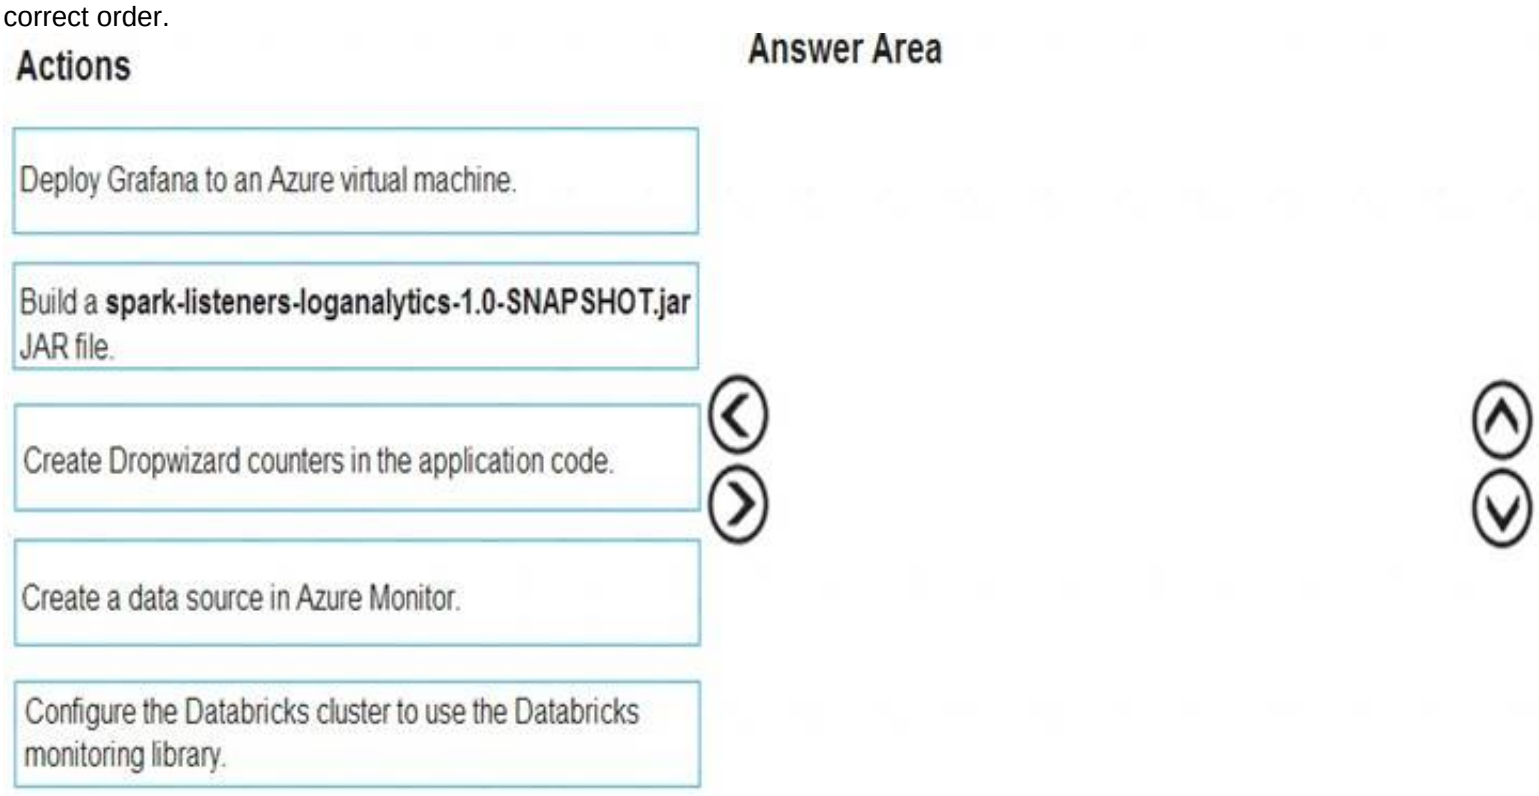

A. Mastered
B. Not Mastered

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 95
- (Exam Topic 5)
You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:

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

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

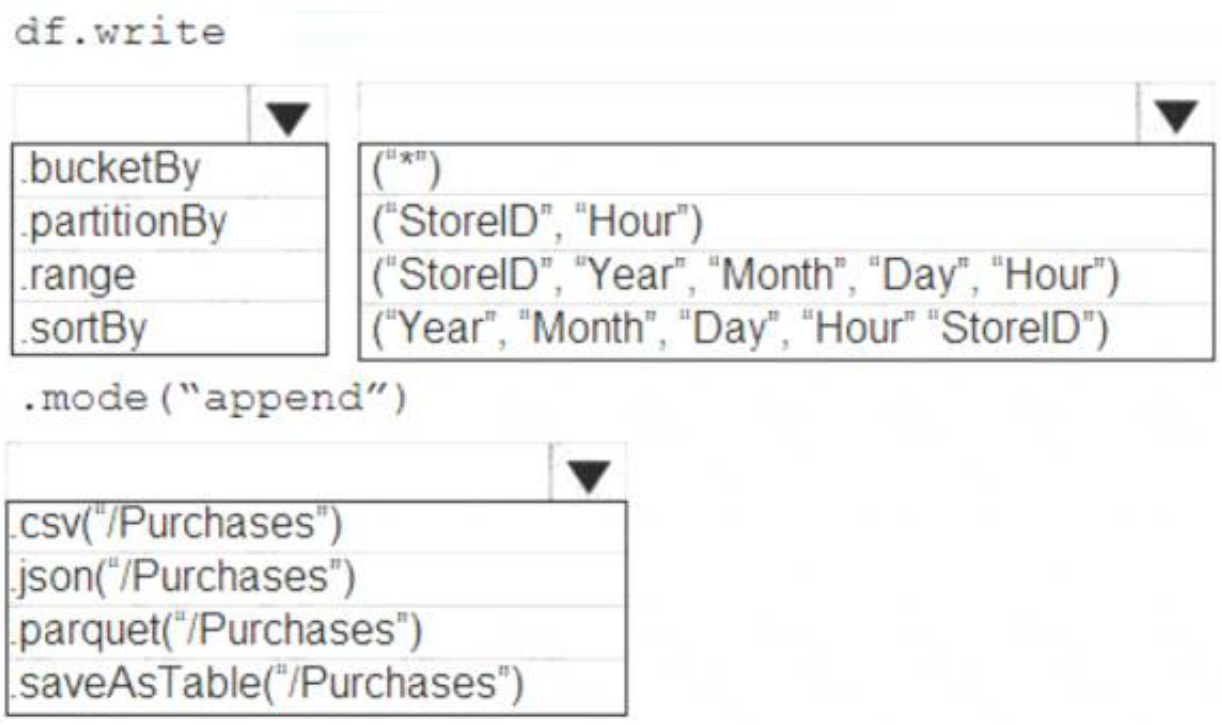

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 99
- (Exam Topic 5)
You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.
You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information.
What should you include in the recommendation?

A. row-level security
B. data masking
C. Always Encrypted
D. column-level security

Answer: B

Explanation:
Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card.
Example:
XXXX-XXXX-XXXX-1234

NEW QUESTION 104
- (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.

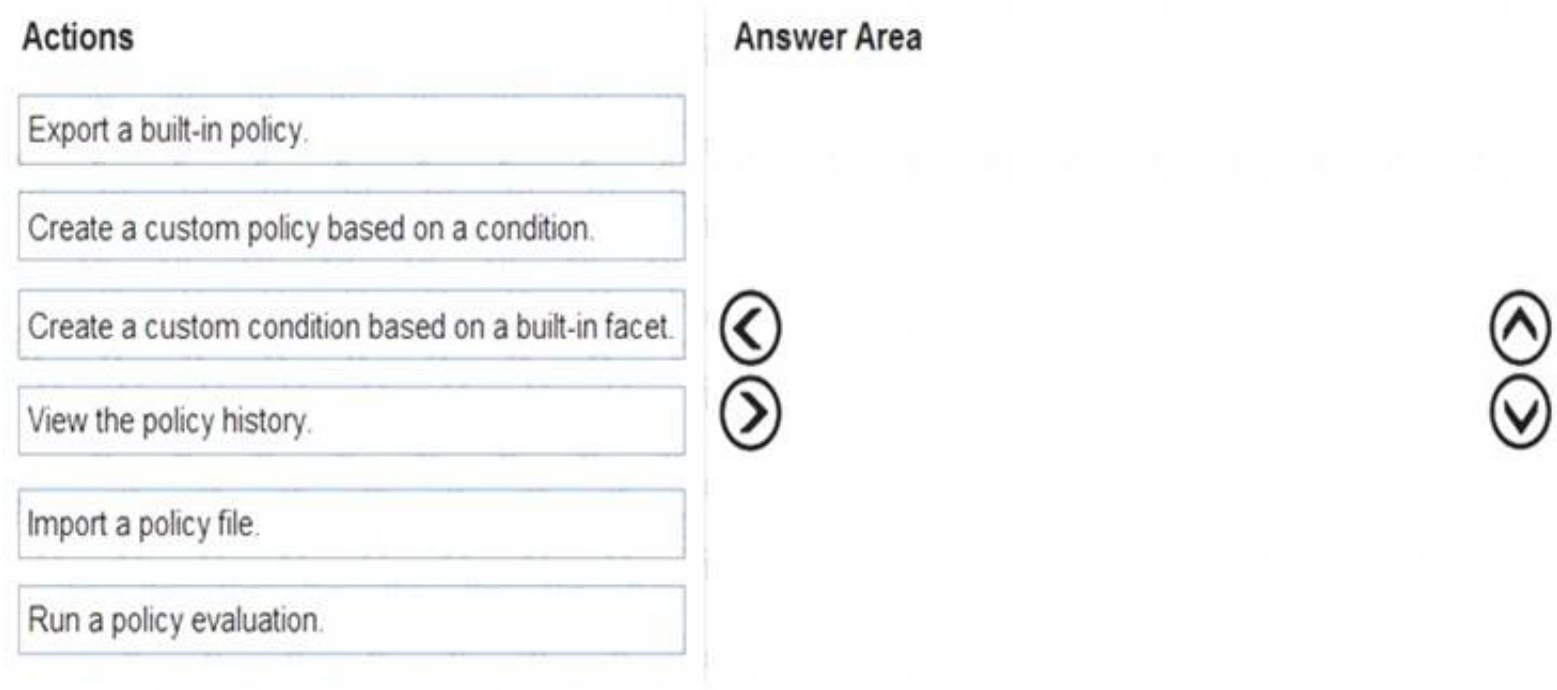

A. Mastered
B. Not Mastered

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 106
- (Exam Topic 5)
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?

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

Answer: A

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

NEW QUESTION 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 SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance.
Solution: You change the data file for the master database to autogrow by 10 percent. Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Reference:

https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

NEW QUESTION 111
- (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.

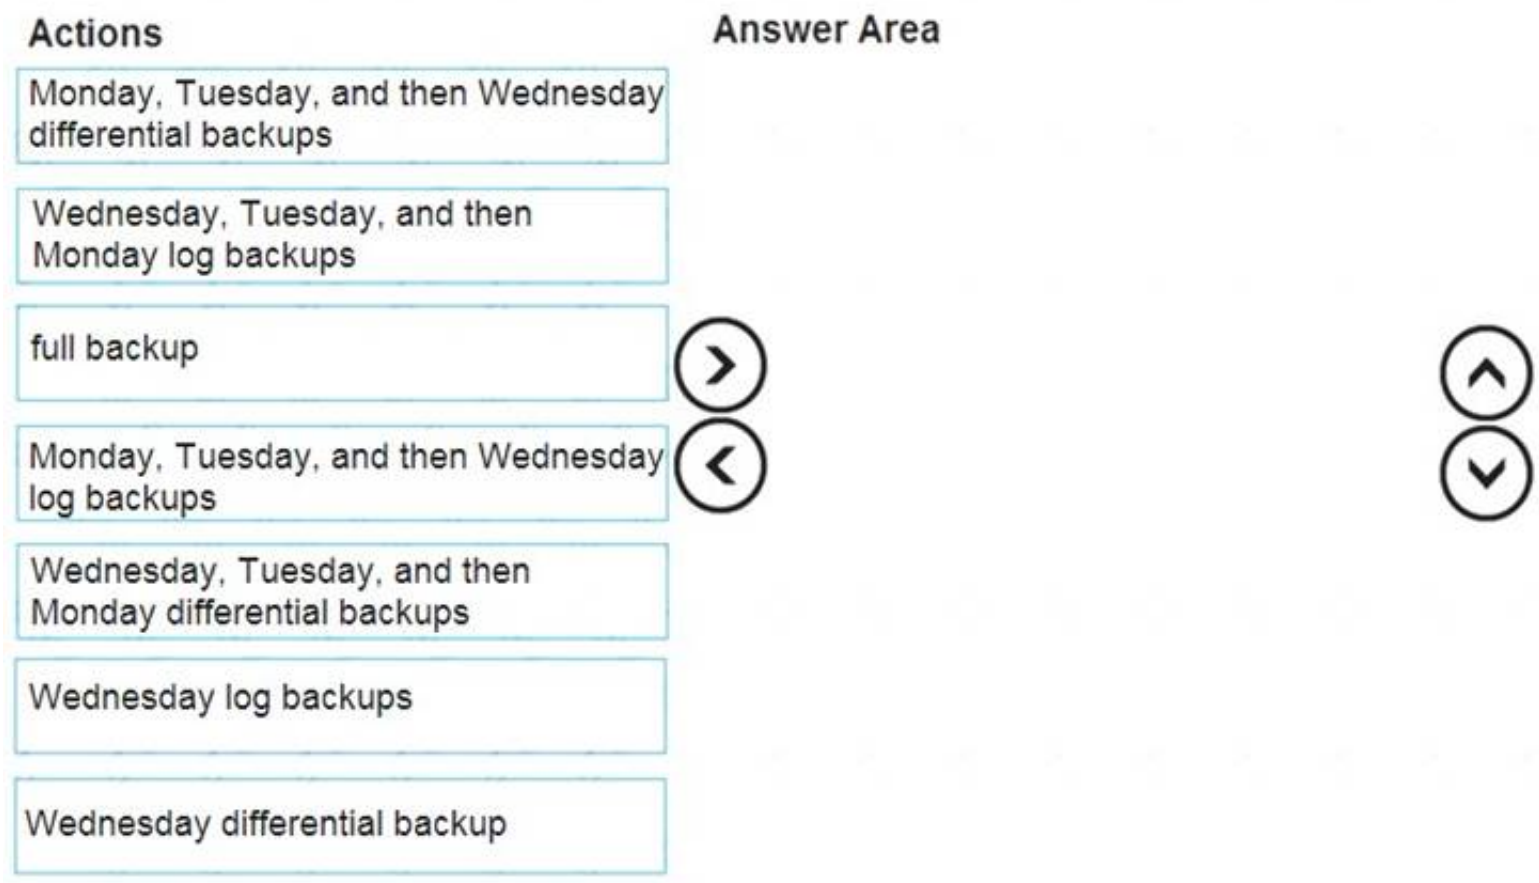

A. Mastered
B. Not Mastered

Answer: A

Explanation:

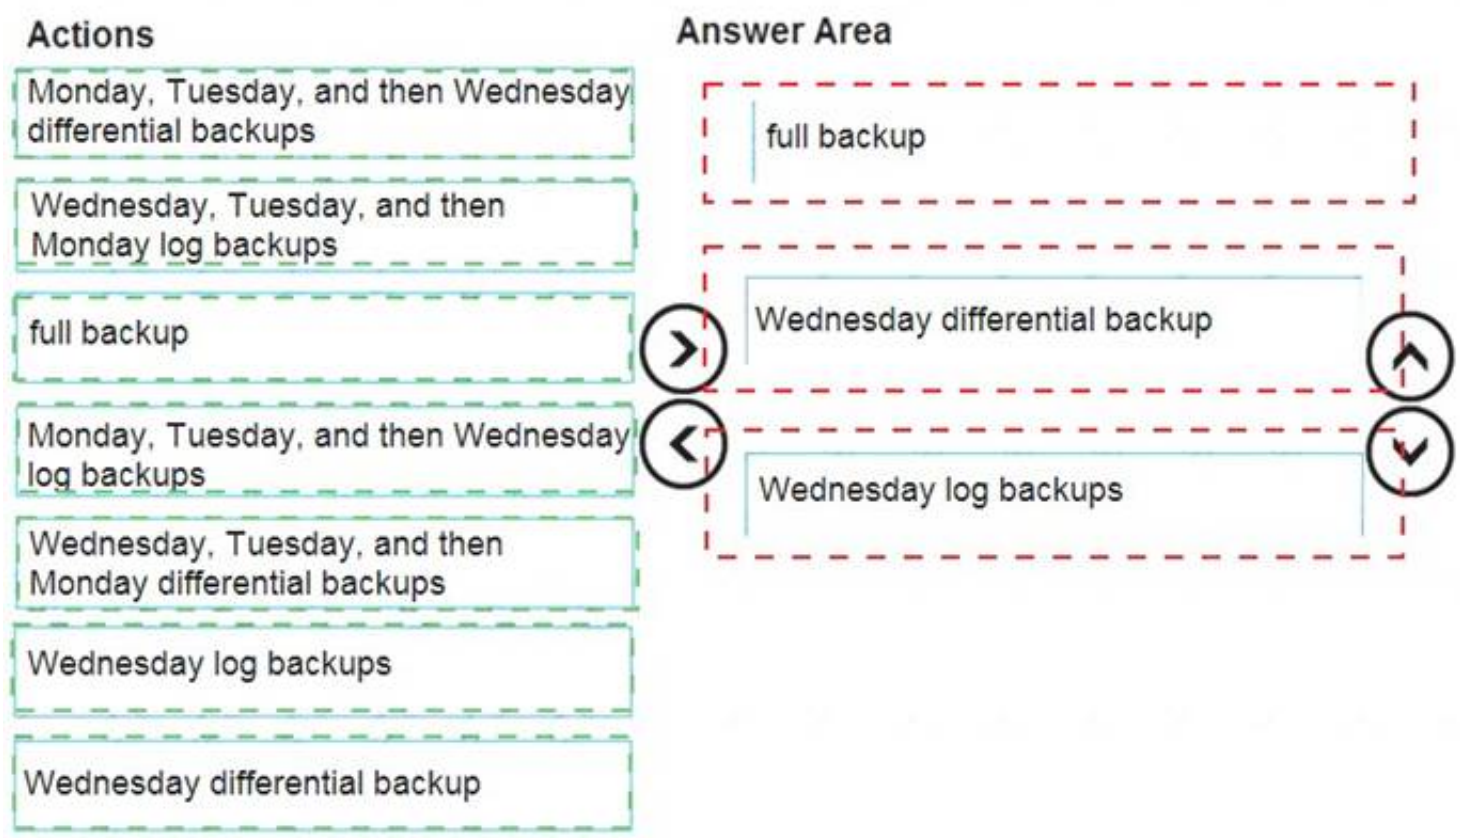

NEW QUESTION 114
- (Exam Topic 5)
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

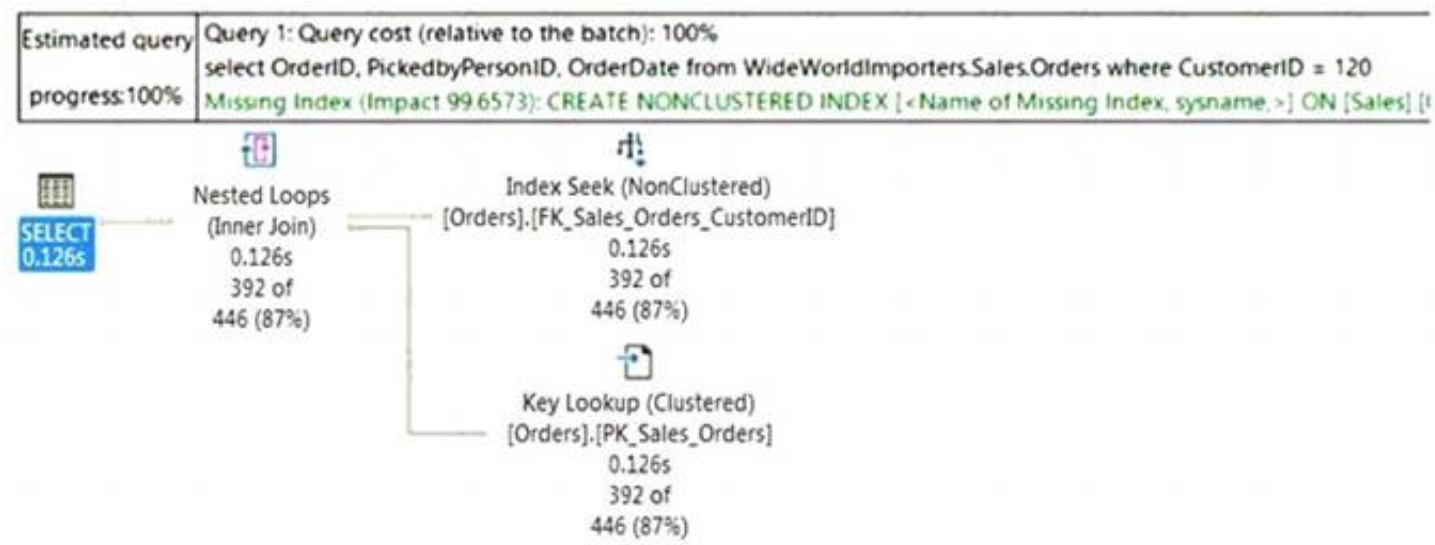

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

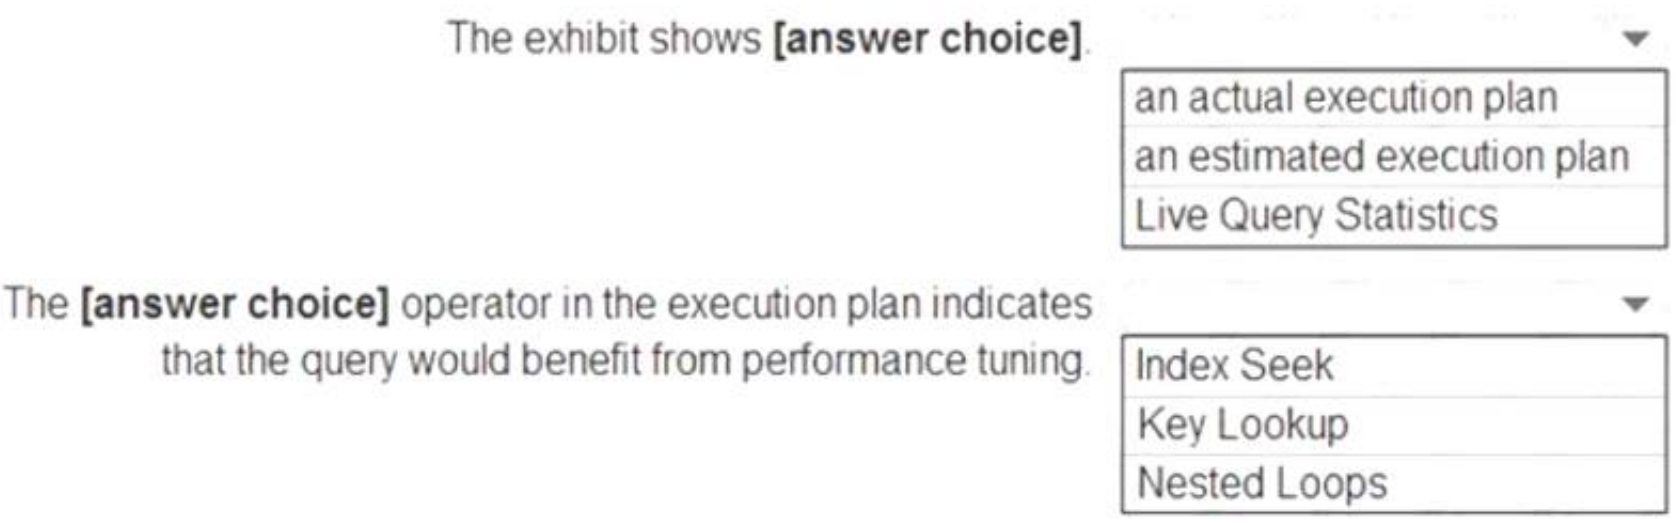

A. Mastered
B. Not Mastered

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 116
- (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 121
- (Exam Topic 5)
You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.
Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.
You set the Read/Write failover policy to Automatic.
What are two results of the configuration? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Answer: AE

Explanation:
A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.
E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be

set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

NEW QUESTION 126
- (Exam Topic 5)
You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. The solution must minimize downtime.
What should you do?

A. Configure Transaction Log Shipping.
B. Implement Always On availability groups.
C. Configure transactional replication.
D. Import a BACPAC.

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1-migra

NEW QUESTION 131
- (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 136
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1.
You need to enable Resource Governor for a App1. The solution must meet the following requirements: App1 must be able to consume all available CPU resources.
App1 must have at least half of the available CPU resources always available.
Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

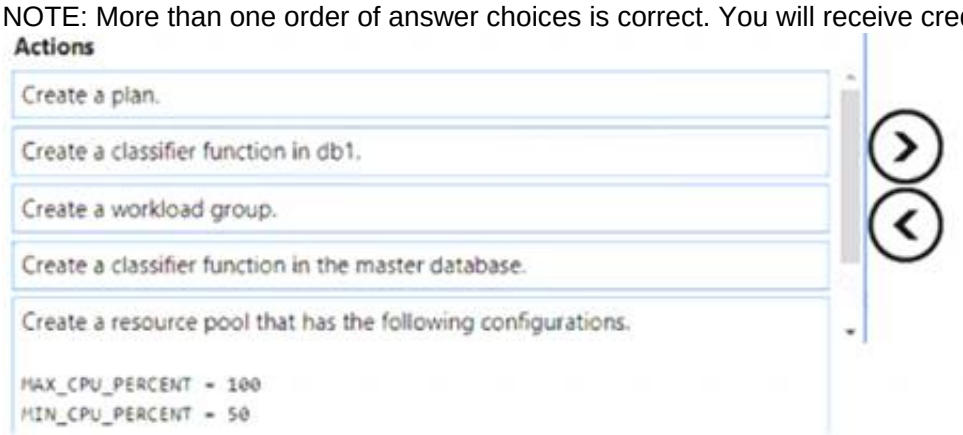

A. Mastered
B. Not Mastered

Answer: A

Explanation:

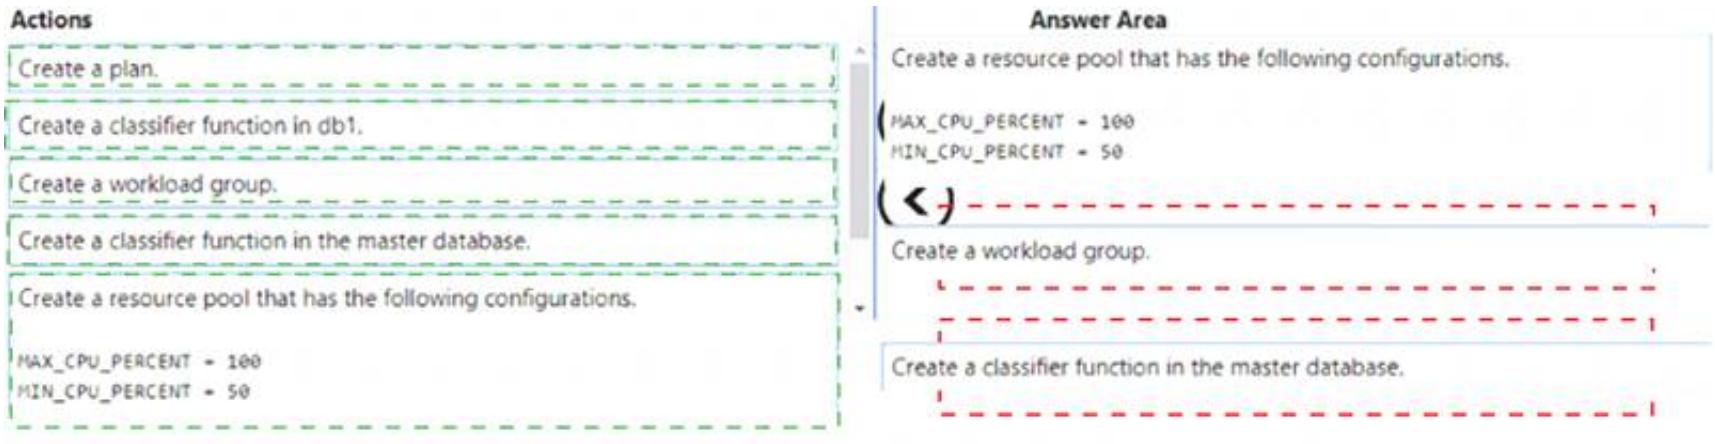

NEW QUESTION 139
- (Exam Topic 5)
You have an Azure SQL database.
Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.
You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.
What should you use?

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

Answer: C

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

NEW QUESTION 142
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.
You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.
The on-premises network connects to VNET1 by using an ExpressRoute connection.
You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service. You need to configure VNET1 to support the migration.
What should you do?

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

Answer: A

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

NEW QUESTION 145
- (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 149
- (Exam Topic 5)

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

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

Answer: D

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

NEW QUESTION 151
- (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 155
- (Exam Topic 5)
You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.
You execute the Transact-SQL query shown in the following exhibit.

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

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

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

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

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

- a header
- a data
- an empty

A. Mastered
B. Not Mastered

Answer: A

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

https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset

NEW QUESTION 157
- (Exam Topic 5)
You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.
Which switch should you use to switch between languages?

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

Answer: B

Explanation:

You can override the default language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.
Reference:
https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use

NEW QUESTION 162
- (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 167
- (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 171
- (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 172
- (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 177
- (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.

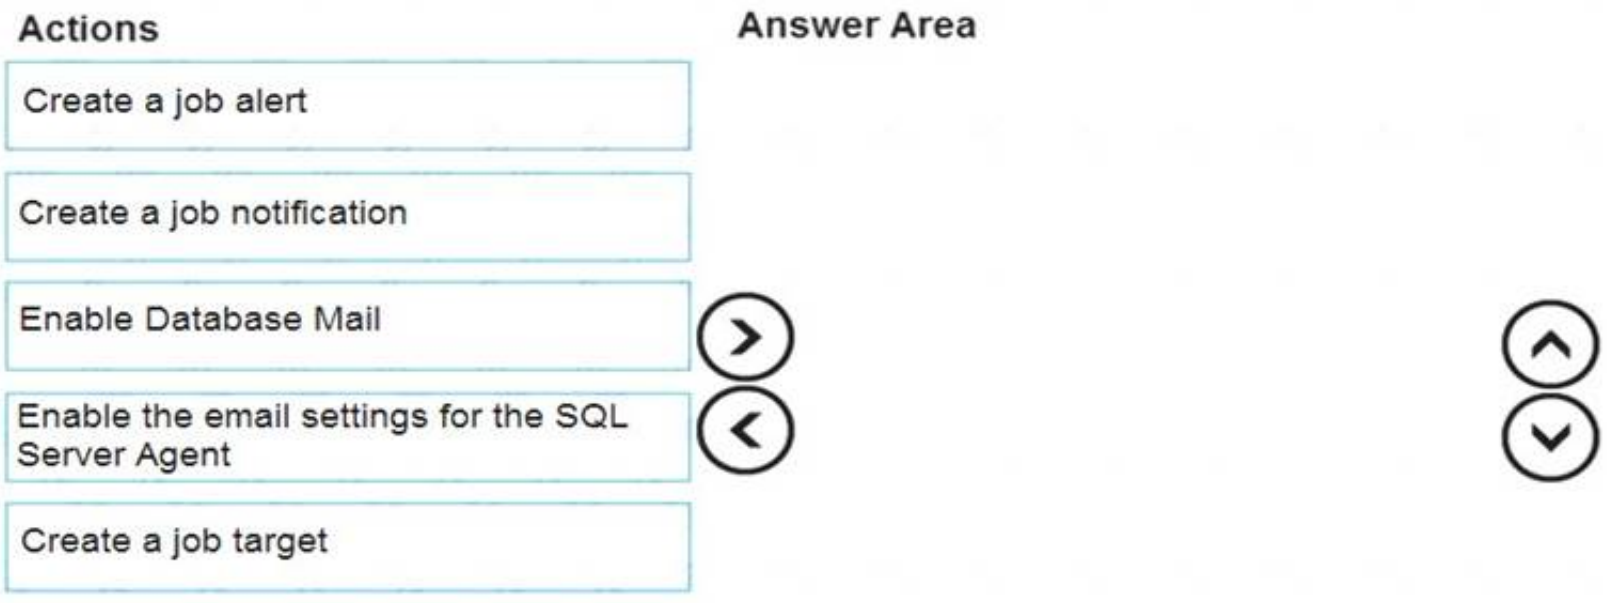

A. Mastered
B. Not Mastered

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 178
- (Exam Topic 5)
You have an Azure SQL Database managed instance. The instance starts experiencing performance issues.
You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.
What should you use?

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

Answer: C

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

NEW QUESTION 179
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance. Solution: You create additional tempdb files. Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

NEW QUESTION 182
- (Exam Topic 5)
You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2. You need to limit the number of IOPs that App2 queries generate on SQL1.
Which two actions should you perform on SQL1? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Answer: BC

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

NEW QUESTION 186
- (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 187
- (Exam Topic 5)
You have an Azure subscription that contains the resources shown in the following table.

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

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

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

Answer: B

NEW QUESTION 188
- (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 191
- (Exam Topic 5)
You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?

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

Answer: D

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

NEW QUESTION 195
- (Exam Topic 5)
You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.
You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

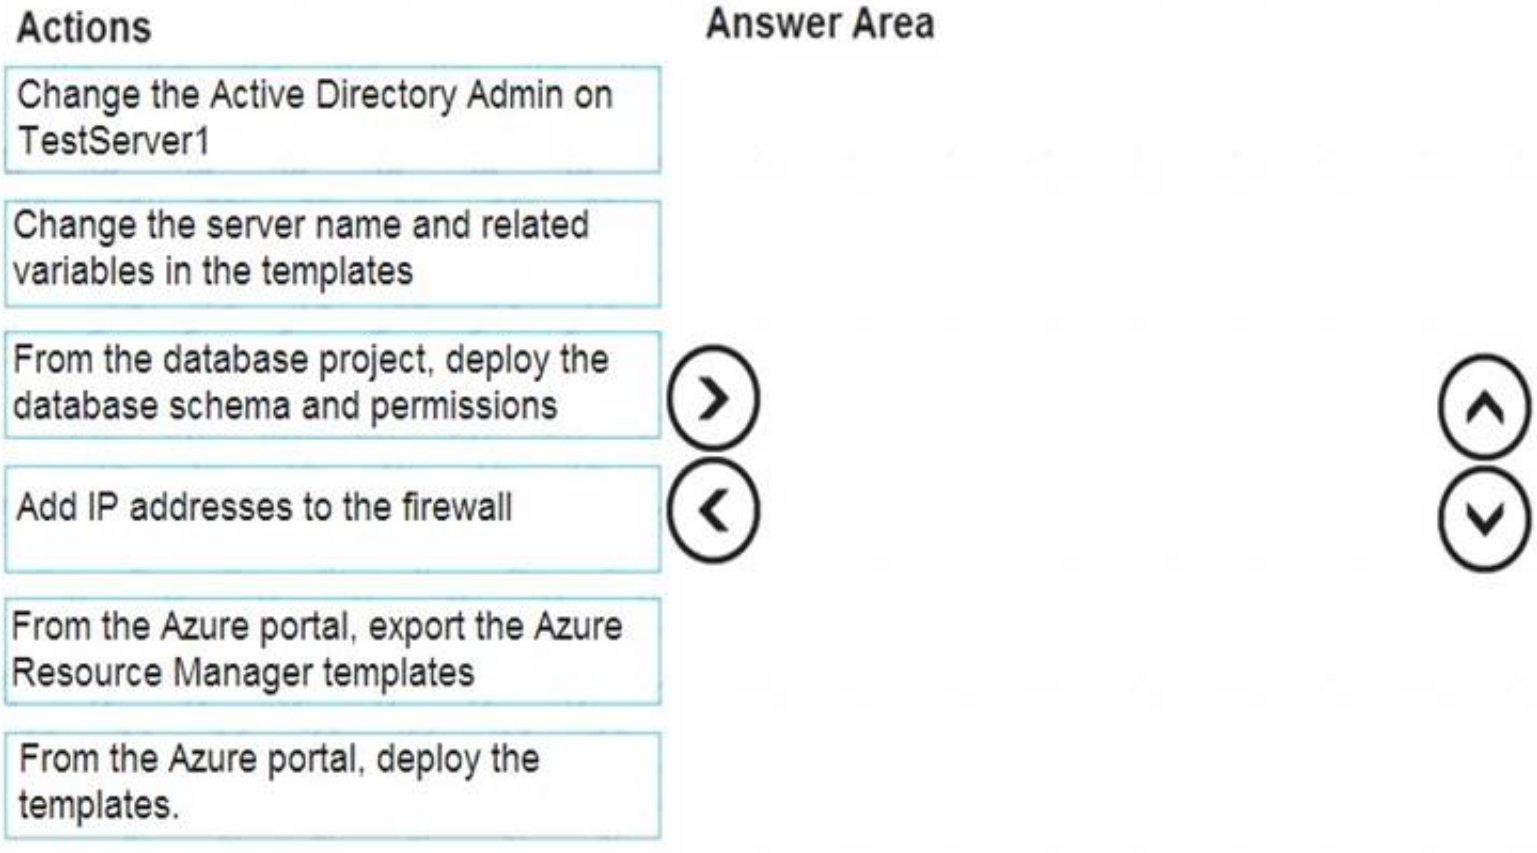

A. Mastered
B. Not Mastered

Answer: A

Explanation:

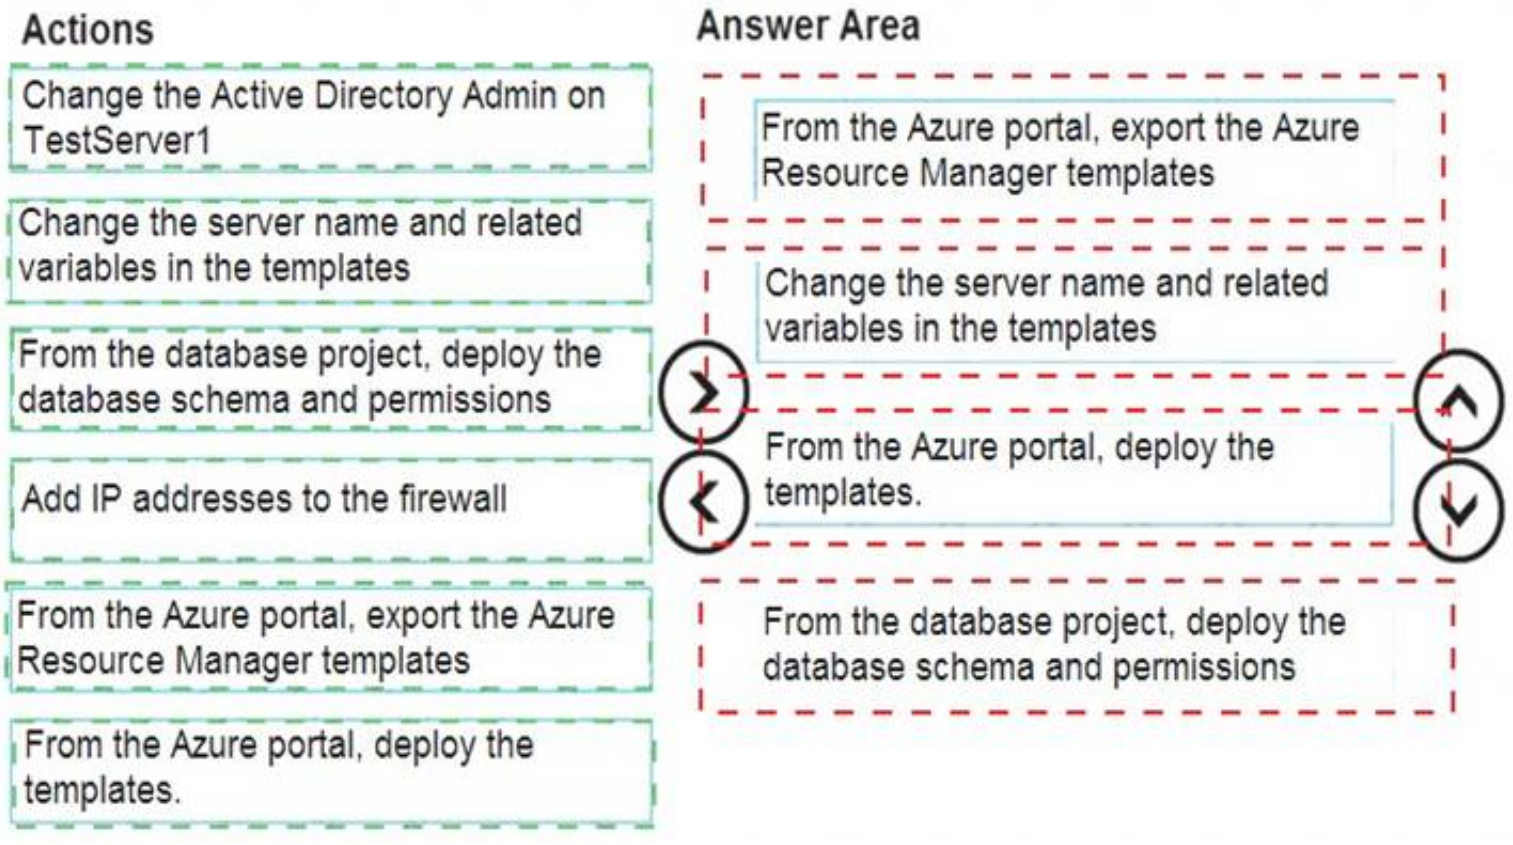

NEW QUESTION 196
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have 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 200
- (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.

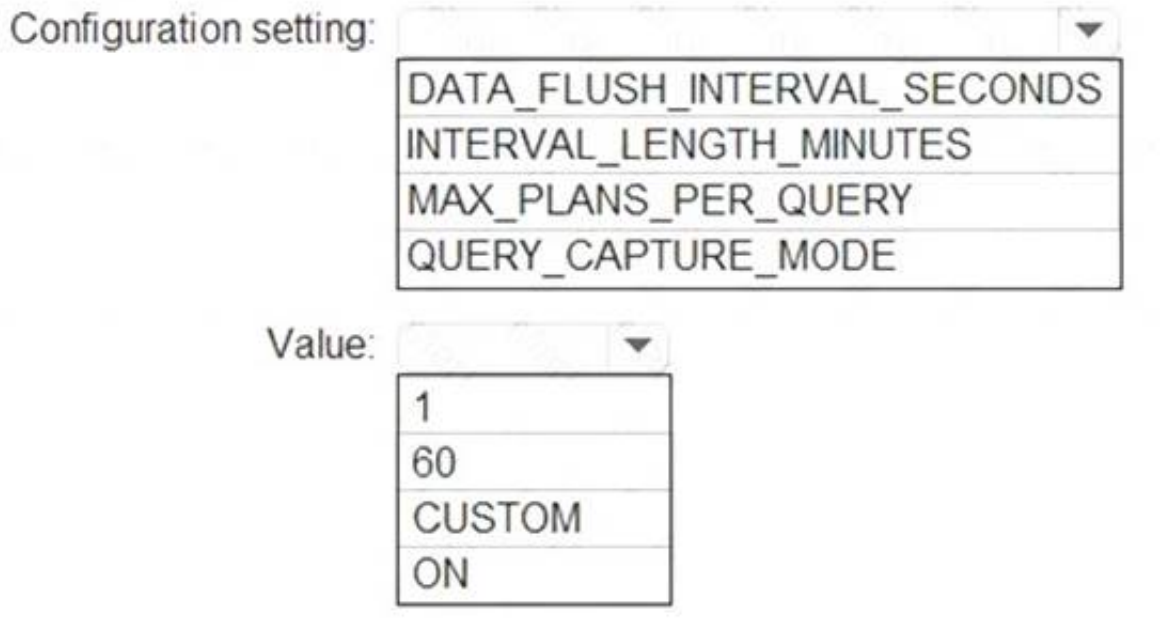

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 205
- (Exam Topic 5)
You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.
You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

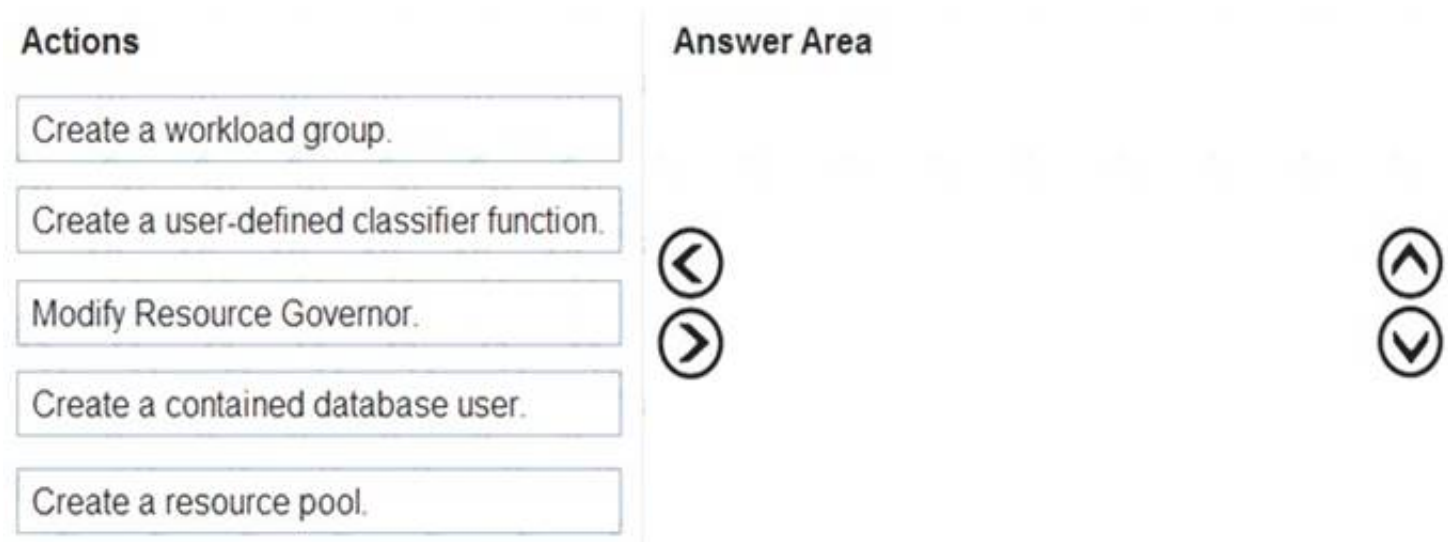

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 208
- (Exam Topic 5)
You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.
You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

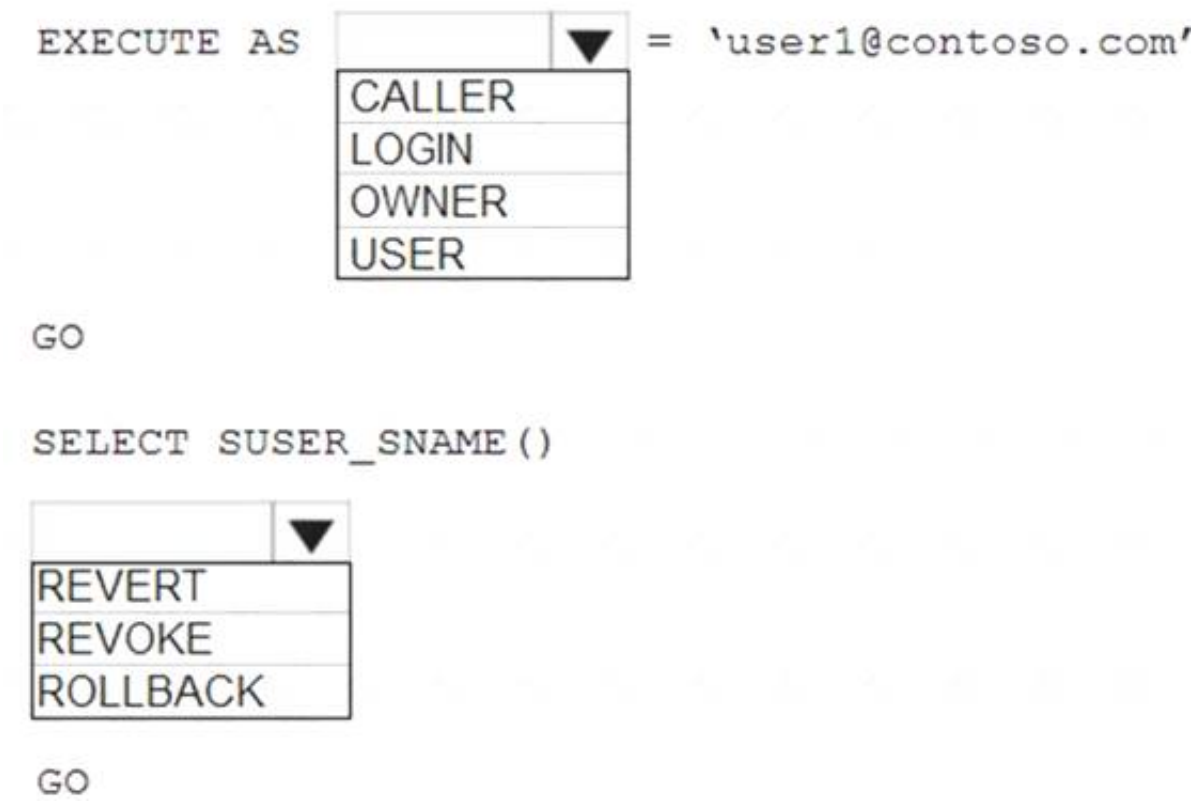

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 212
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named Db1.
You need to enable automatic tuning for Db1.
How should you complete the statements? To answer, select the appropriate answer in the answer area.
NOTE: Each correct selection is worth one point.

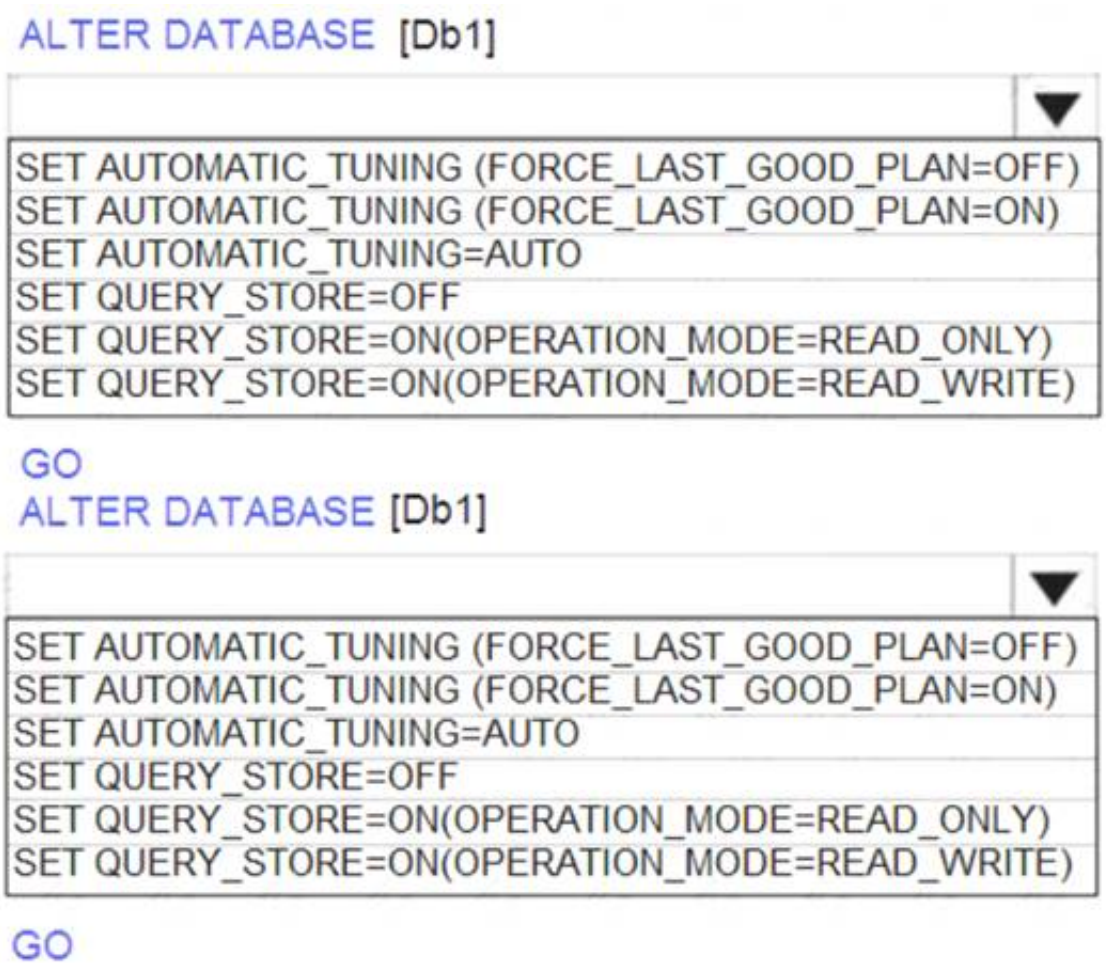

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 213
- (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 214
- (Exam Topic 5)
You have an Azure SQL database.
You have a query and the associated execution plan as shown in the following exhibit.

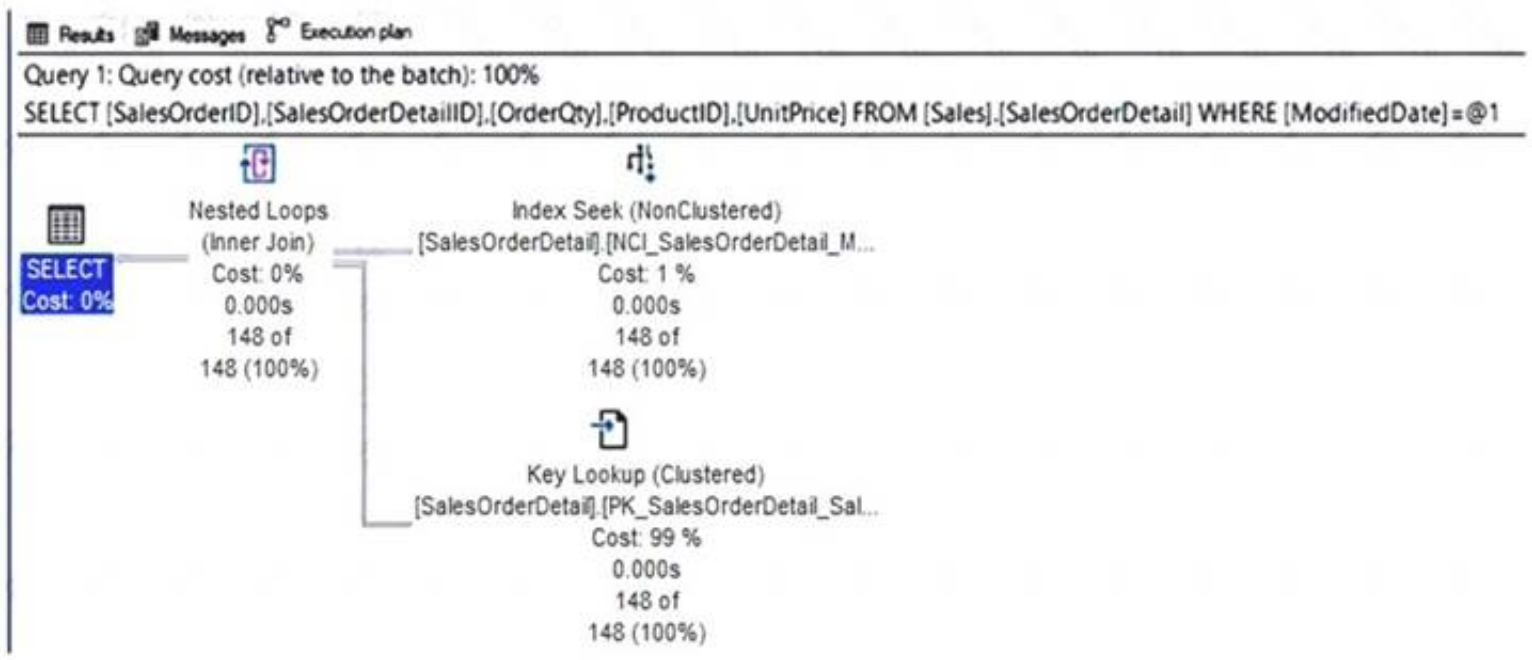

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

The performance issue stems from the **[answer choice]** operator.

- Select
- Index Seek
- Key Lookup
- Nested Loops

The performance issue can be resolved by adding include columns to the **[answer choice]**.

- heap
- clustered index
- nonclustered index

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 215
- (Exam Topic 5)
You have an 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 216
- (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 219
- (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.

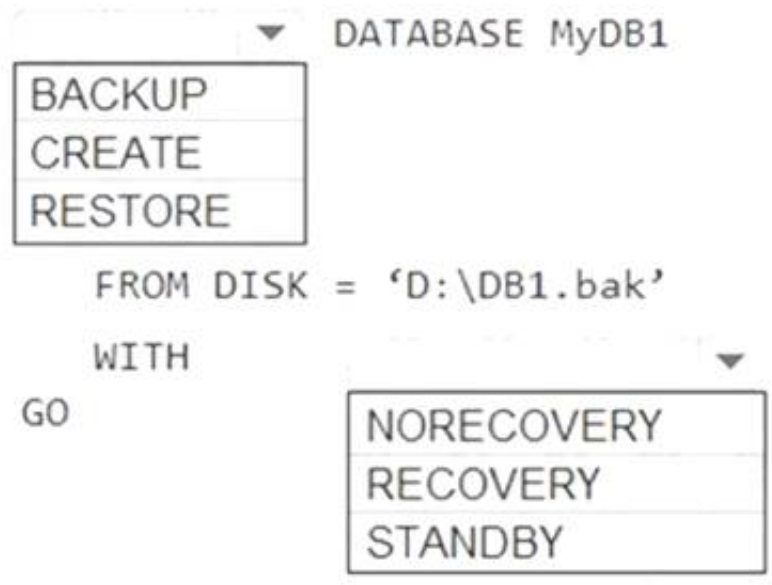

A. Mastered
B. Not Mastered

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 223
- (Exam Topic 5)
You have a data warehouse in Azure Synapse Analytics.
You need to ensure that the data in the data warehouse is encrypted at rest. What should you enable?

A. Transparent Data Encryption (TDE)
B. Advanced Data Security for this database
C. Always Encrypted for all columns

D. Secure transfer required

Answer: A

Explanation:
Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview

NEW QUESTION 224
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine.
You review the query plan shown in the following exhibit.

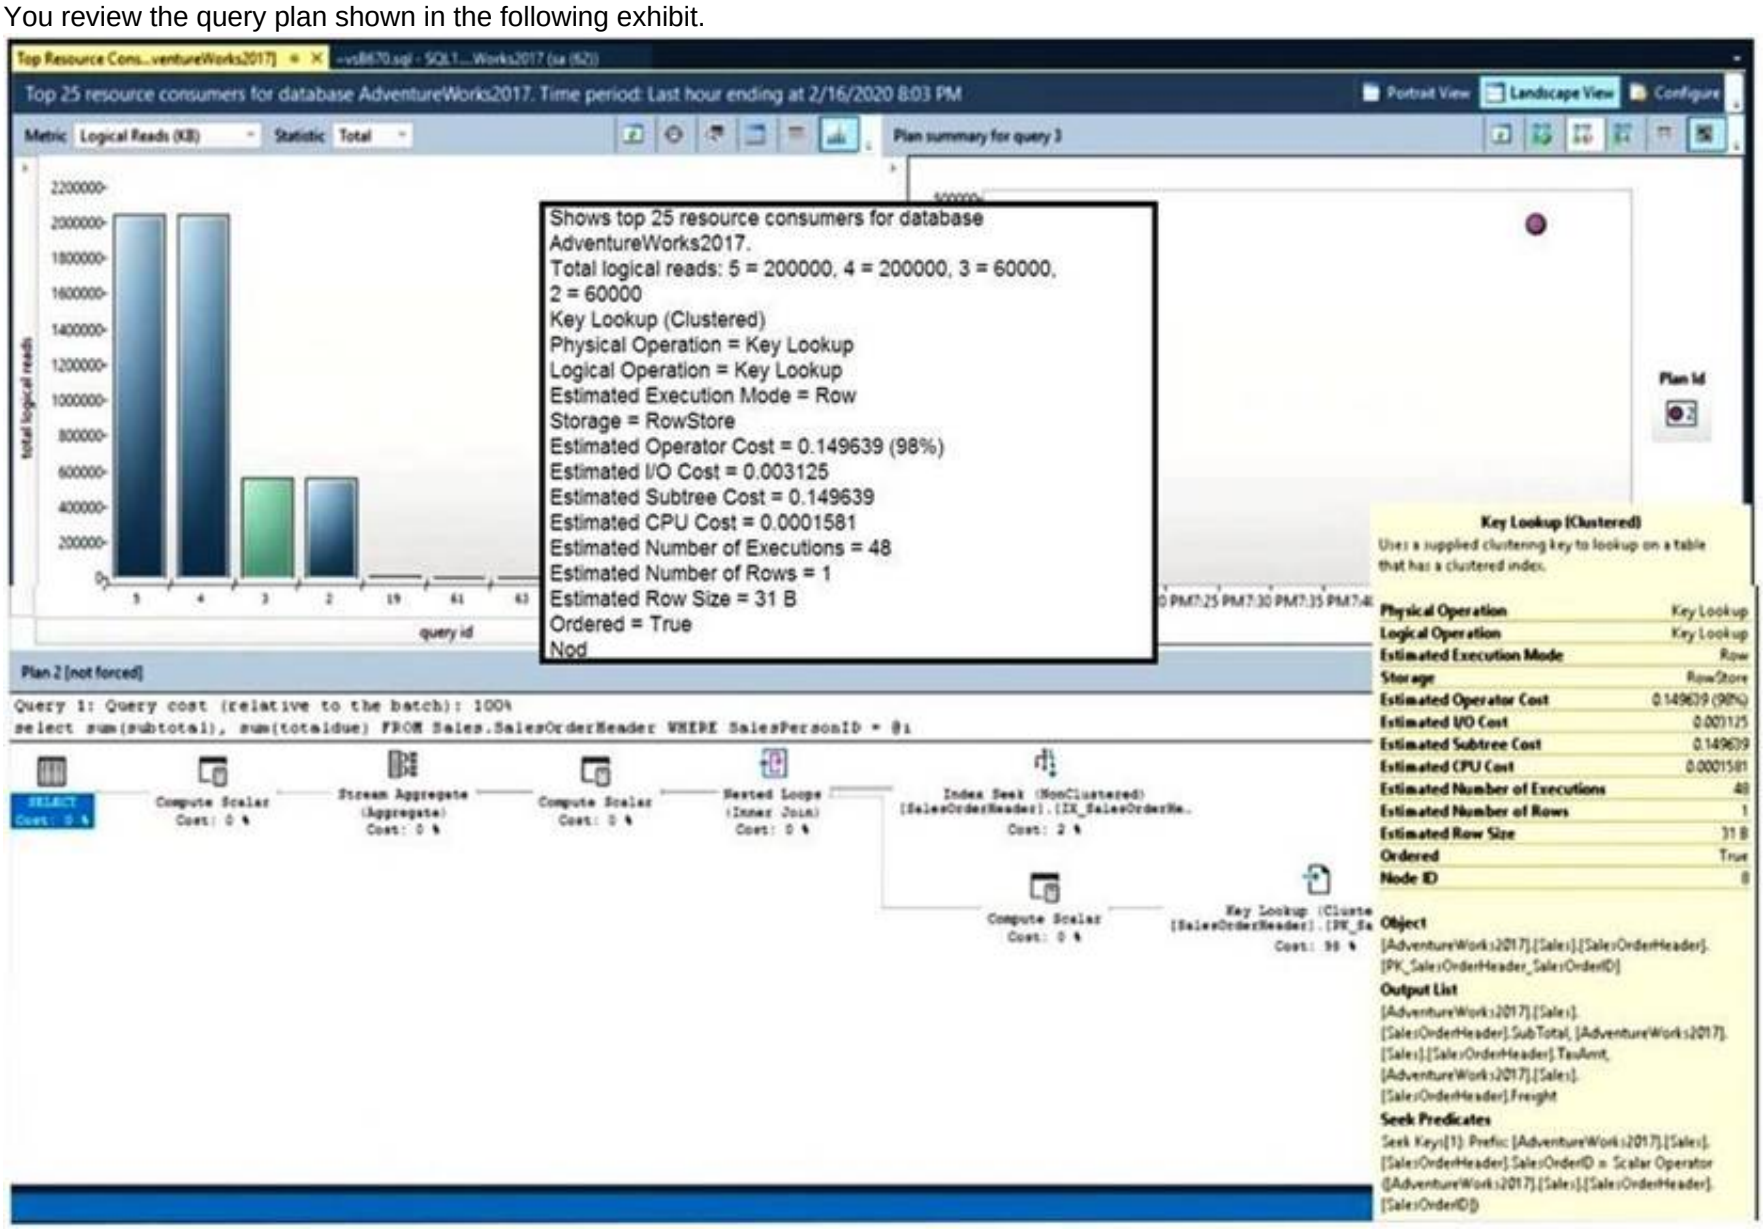

For each of the following statements, select yes if the statement is true. Otherwise, select no.
NOTE: Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 226
- (Exam Topic 5)
You have an Azure SQL database named DB3.
You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click the Exhibit tab.)

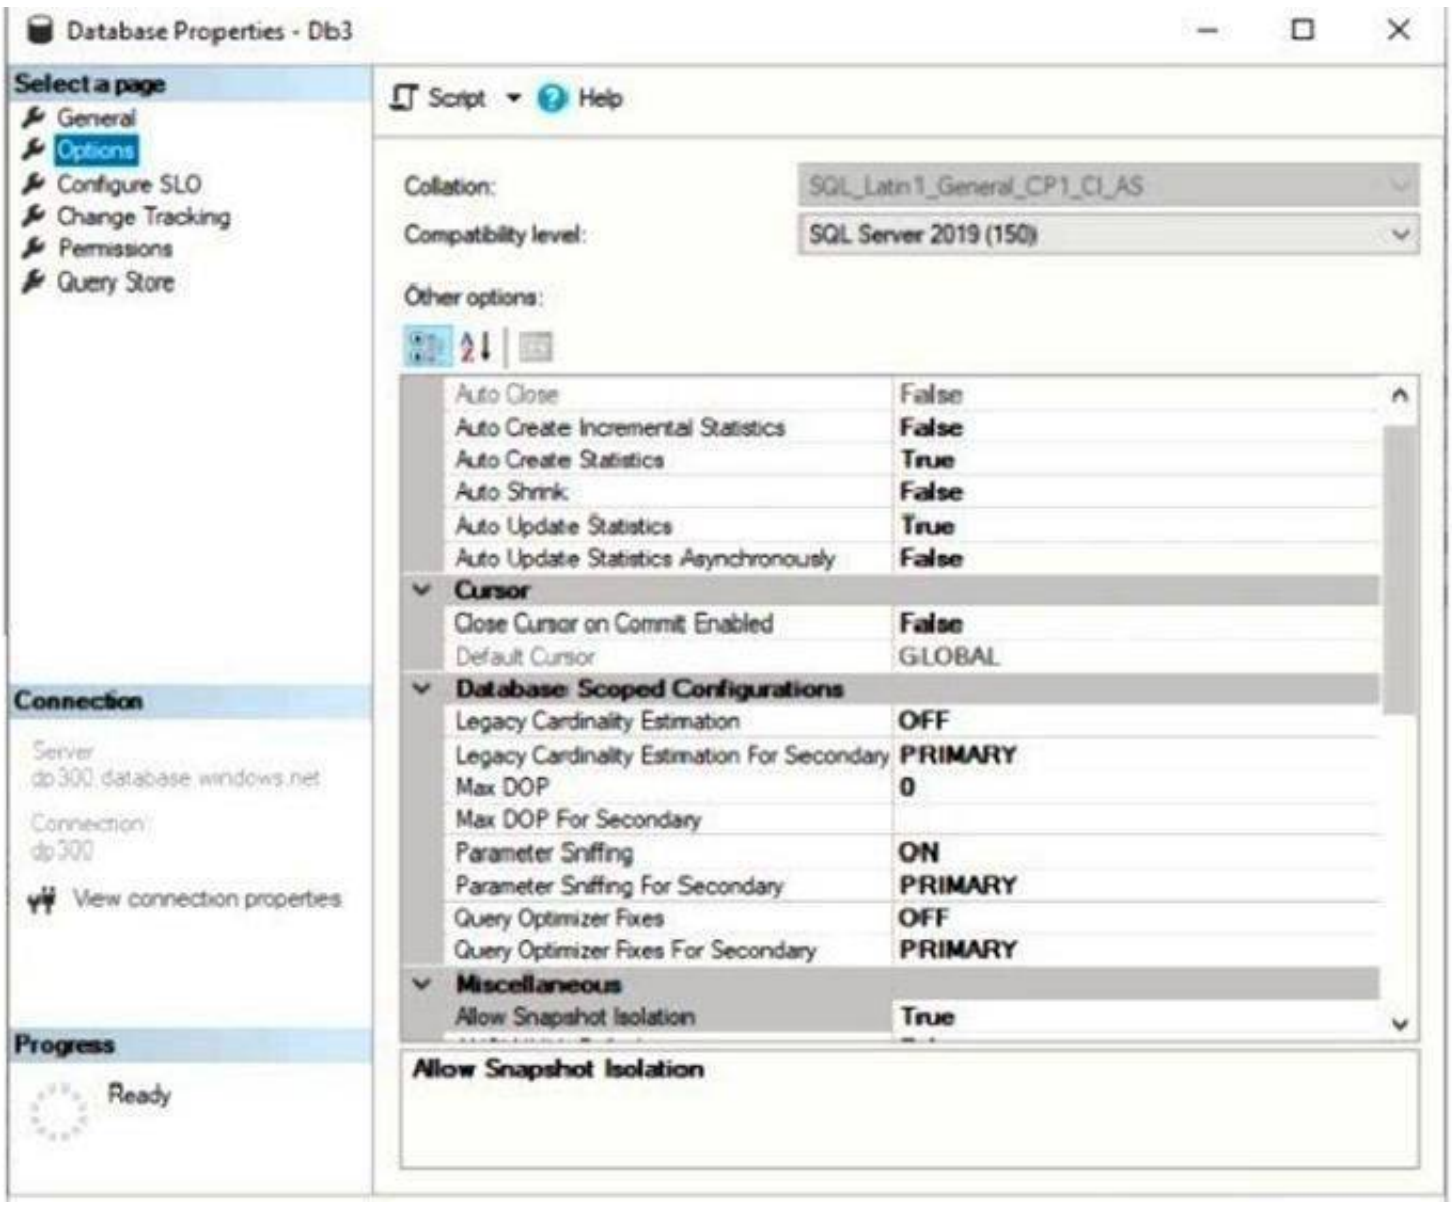

Which Transact-SQL command should you run?

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

Answer: C

Explanation:
The exhibits displays Database [State] properties.
To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-properties-options-page

NEW QUESTION 231
- (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 232
- (Exam Topic 5)
You have an Azure SQL logical server. You run the following script.

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

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

NOTE: Each correct selection is worth one point

Statements	Yes	No
The `Orders` table will allow only rows to be inserted.	○	○
To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.	○	○
To ensure that a timestamp is added to each row in the `Orders` table, the `GENERATED ALWAYS`	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Yes
No No

NEW QUESTION 235
- (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.

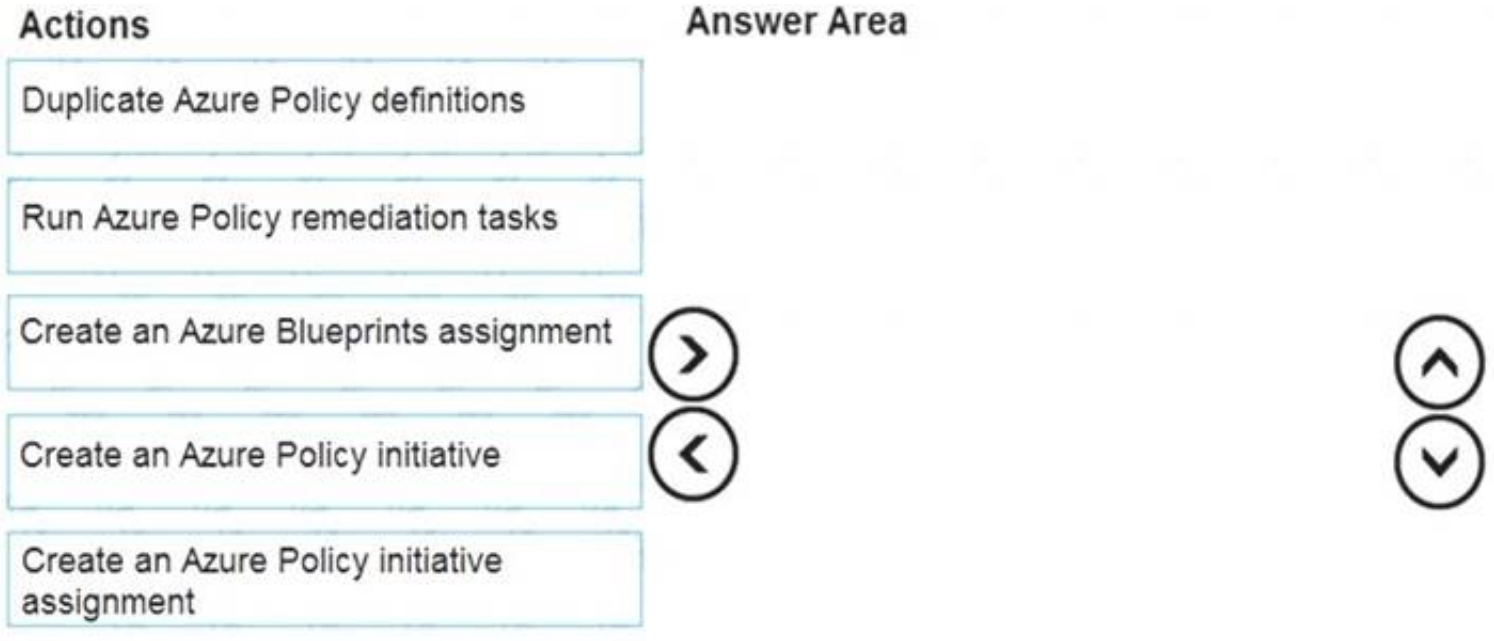

A. Mastered
B. Not Mastered

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 239
- (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 241
- (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 246
- (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 247
- (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 250
- (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 252
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.
Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 254
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1.
Which isolation level should you configure?

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

Answer: B

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

NEW QUESTION 256
- (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 258
- (Exam Topic 4)
You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application Description automatically generated
Box 1: Hash Scenario:
Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.
A hash distributed table can deliver the highest query performance for joins and aggregations on large tables. Box 2: Round-robin
Scenario:
You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.
A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:

- When you cannot identify a single key to distribute your data.
- If your data doesn't frequently join with data from other tables.
- When there are no obvious keys to join.

Reference:
https://rajanieshkaushikk.com/2020/09/09/how-to-choose-right-data-distribution-strategy-for-azure-synapse/

NEW QUESTION 261
- (Exam Topic 4)
You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.
Which Azure Storage functionality should you include in the solution?

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

Answer: C

Explanation:

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

NEW QUESTION 265
- (Exam Topic 3)
Which counter should you monitor for real-time processing to meet the technical requirements?

A. SU% Utilization
B. CPU% utilization
C. Concurrent users
D. Data Conversion Errors

Answer: B

Explanation:
Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.
To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/monitor-tune-overview

NEW QUESTION 270
- (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

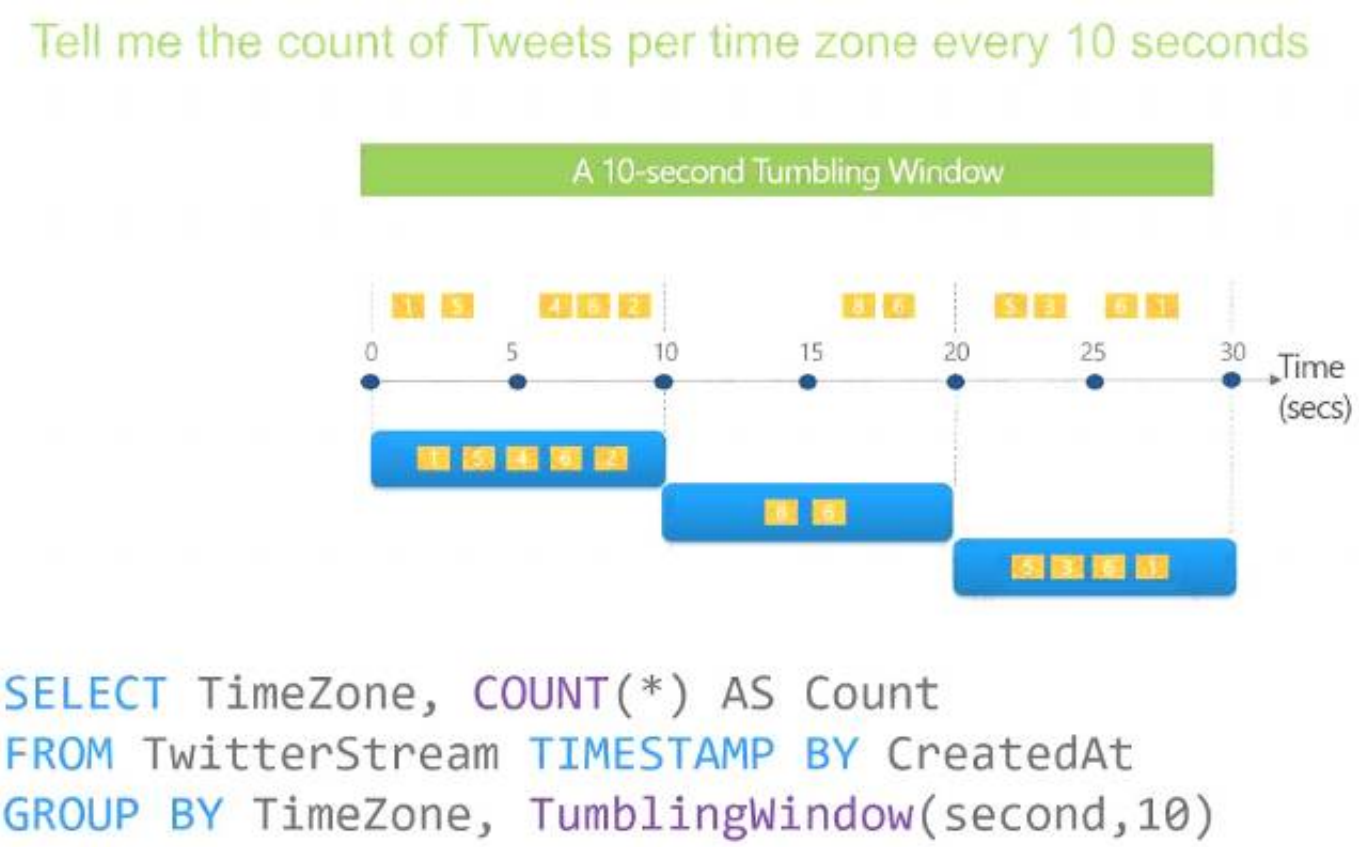

Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 274
- (Exam Topic 2)
You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.
What should you do?

A. Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
B. Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
C. Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
D. Create an Azure Monitor alert rule that has a dynamic threshold and assign the alert rule to an action group.

Answer: D

Explanation:

Reference:
https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/

NEW QUESTION 275
- (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 278
- (Exam Topic 2)
Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

A. Business Critical 4-vCore
B. Hyperscale
C. General Purpose v-vCore
D. Serverless

Answer: A

Explanation:
There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed. Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

NEW QUESTION 282
- (Exam Topic 2)
What should you implement to meet the disaster recovery requirements for the PaaS solution?

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

Answer: B

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

NEW QUESTION 285
- (Exam Topic 2)
Which audit log destination should you use to meet the monitoring requirements?

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

Answer: C

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

NEW QUESTION 288
- (Exam Topic 1)
You are evaluating the business goals.
Which feature should you use to provide customers with the required level of access based on their service agreement?

A. dynamic data masking
B. Conditional Access in Azure
C. service principals
D. row-level security (RLS)

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15

NEW QUESTION 289
- (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.

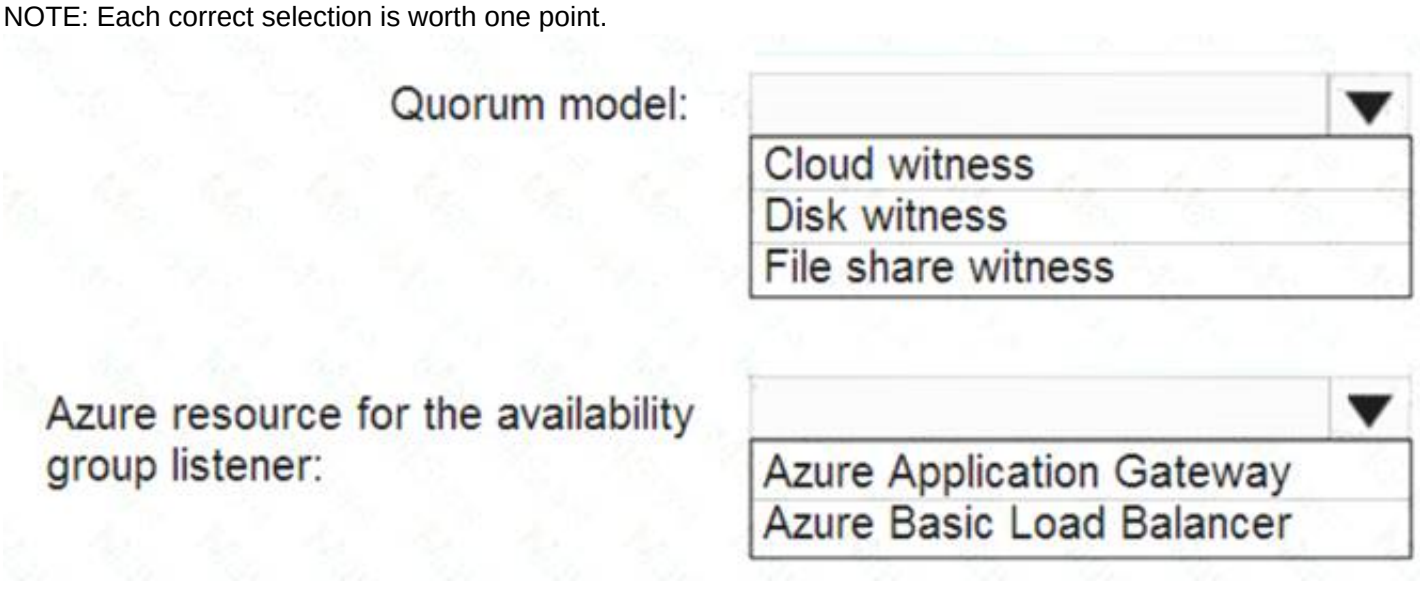

A. Mastered
B. Not Mastered

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 public Virtual 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 292
- (Exam Topic 1)

You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

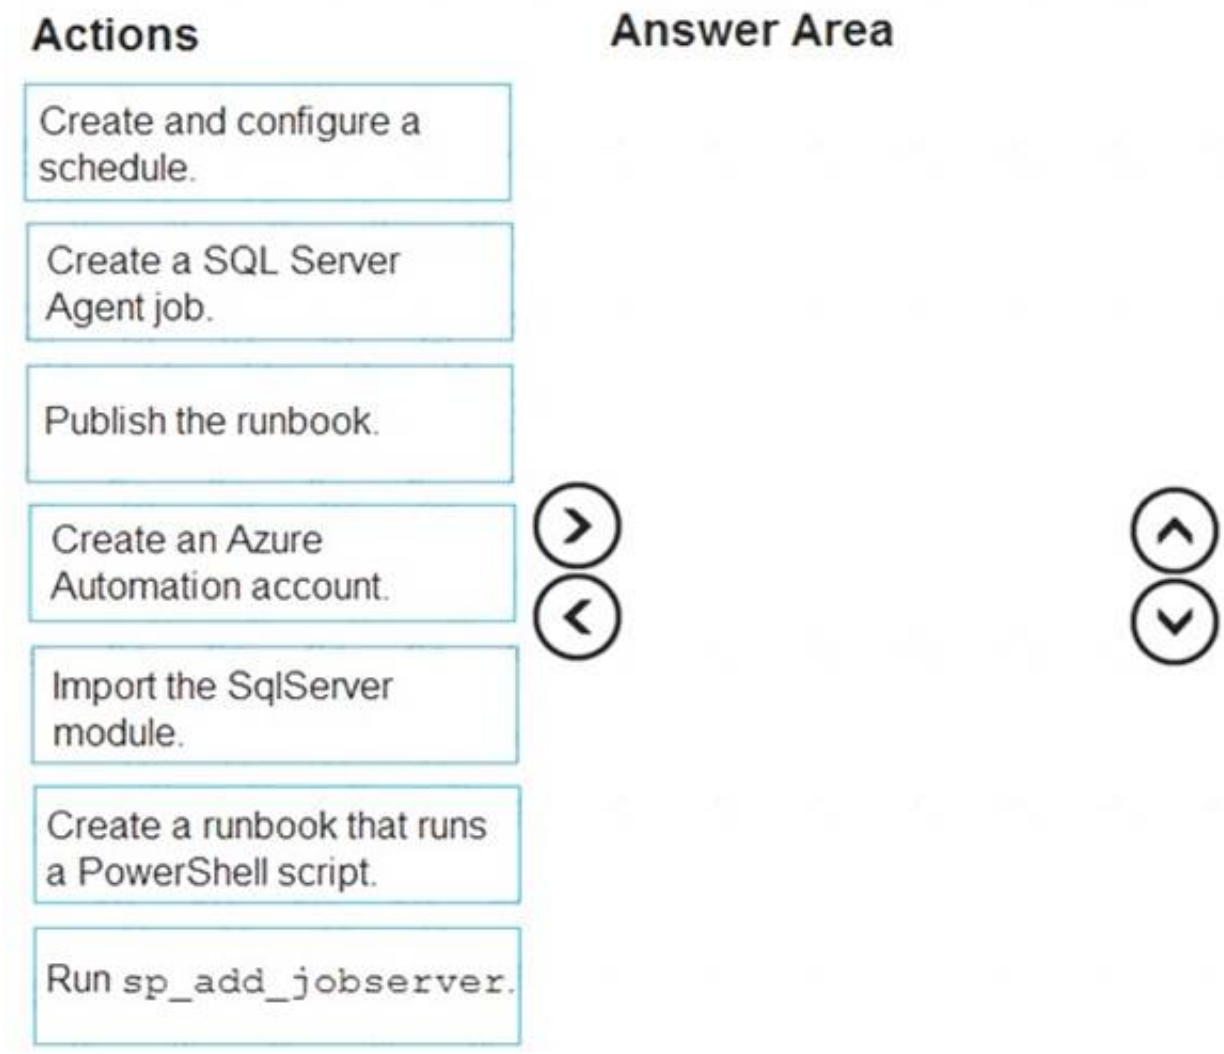

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Automating Azure SQL DB index and statistics maintenance using Azure Automation:
* 1. Create Azure automation account (Step 1)
* 2. Import SQLServer module (Step 2)
* 3. Add Credentials to access SQL DB
This will use secure way to hold login name and password that will be used to access Azure SQL DB
* 4. Add a runbook to run the maintenance (Step 3)
Steps: * 1. Click on "runbooks" at the left panel and then click "add a runbook"
* 2. Choose "create a new runbook" and then give it a name and choose "Powershell" as the type of the runbook and then click on "create"

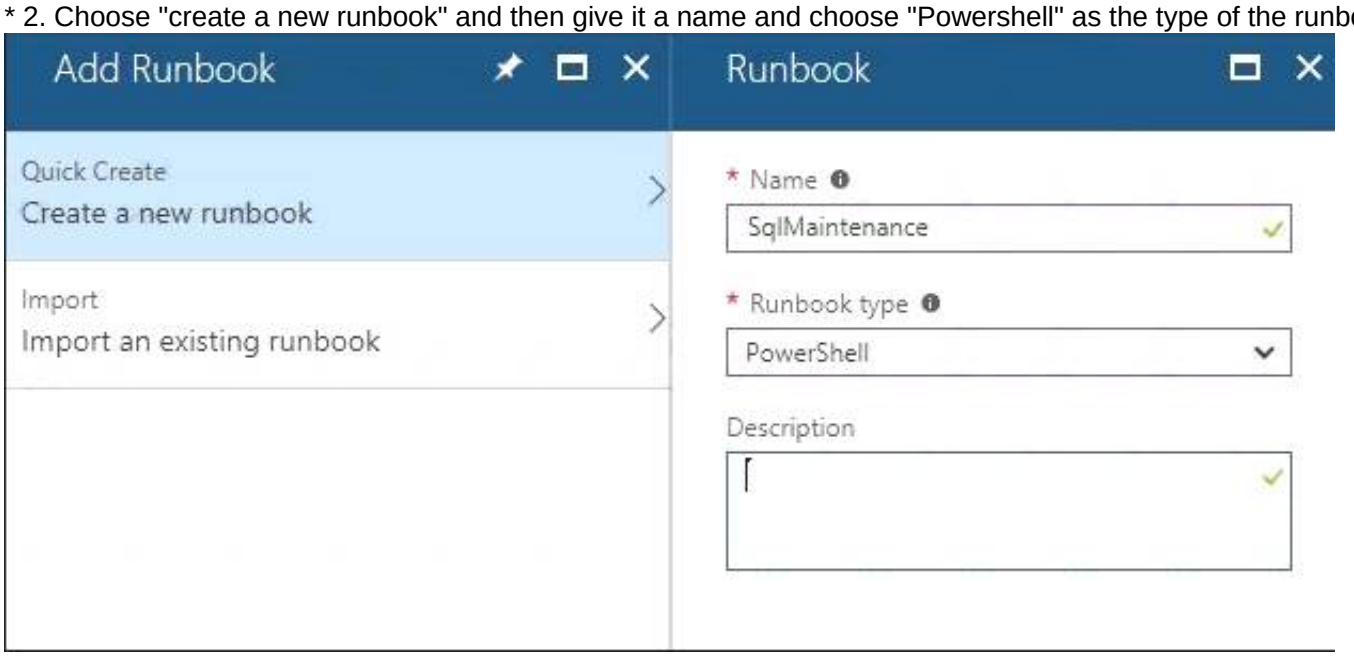

* 5. Schedule task (Step 4)
Steps:1. Click on Schedules2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.
Reference:
https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statist

NEW QUESTION 296
- (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.

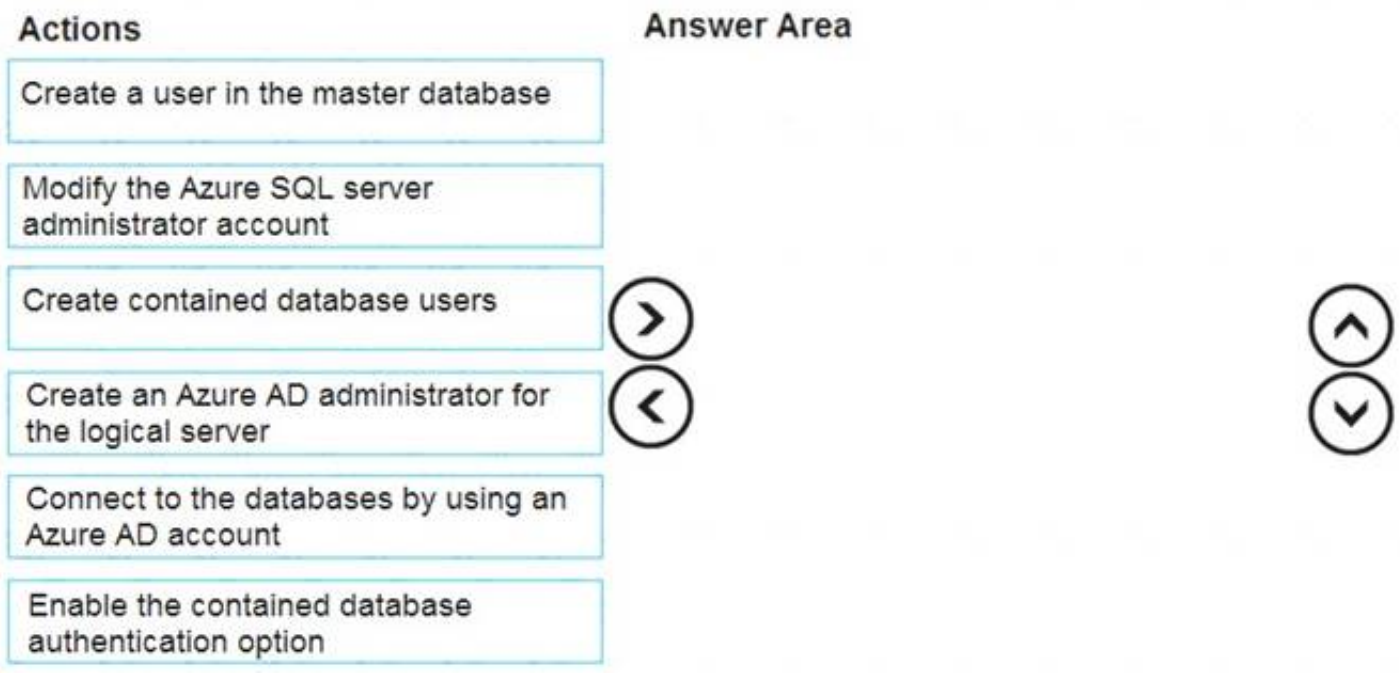

A. Mastered
B. Not Mastered

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 300
.....

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