DP-300 Dumps

Administering Relational Databases on Microsoft Azure (beta)

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

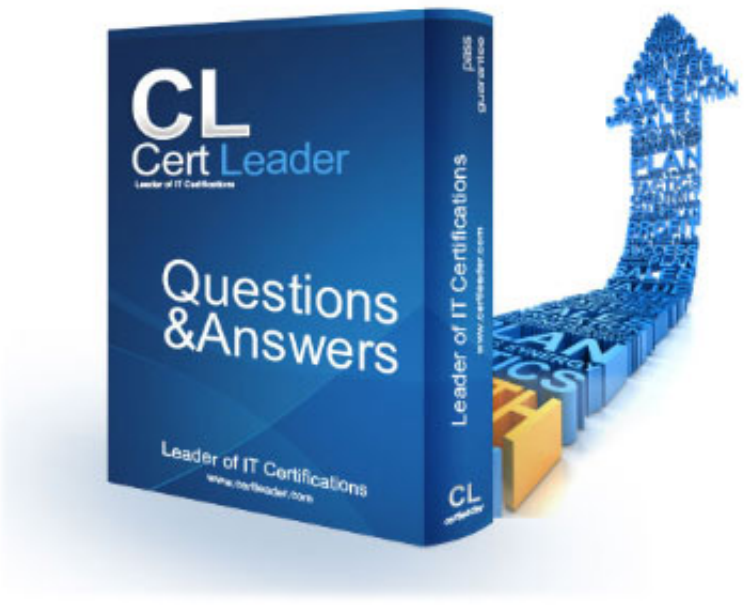

NEW QUESTION 1
- (Exam Topic 5)
You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 2
- (Exam Topic 5)
You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:
- Send the output to an Azure Synapse.
- Identify spikes and dips in time series data.
- Minimize development and configuration effort.

Which should you include in the solution?

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

Answer: C

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

NEW QUESTION 3
- (Exam Topic 5)
You have an instance of SQL Server on Azure Virtual Machines named SQL1. SQL1 contains a database named DB1.
You need to enable Transparent Data Encryption (TDE) for DB1.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

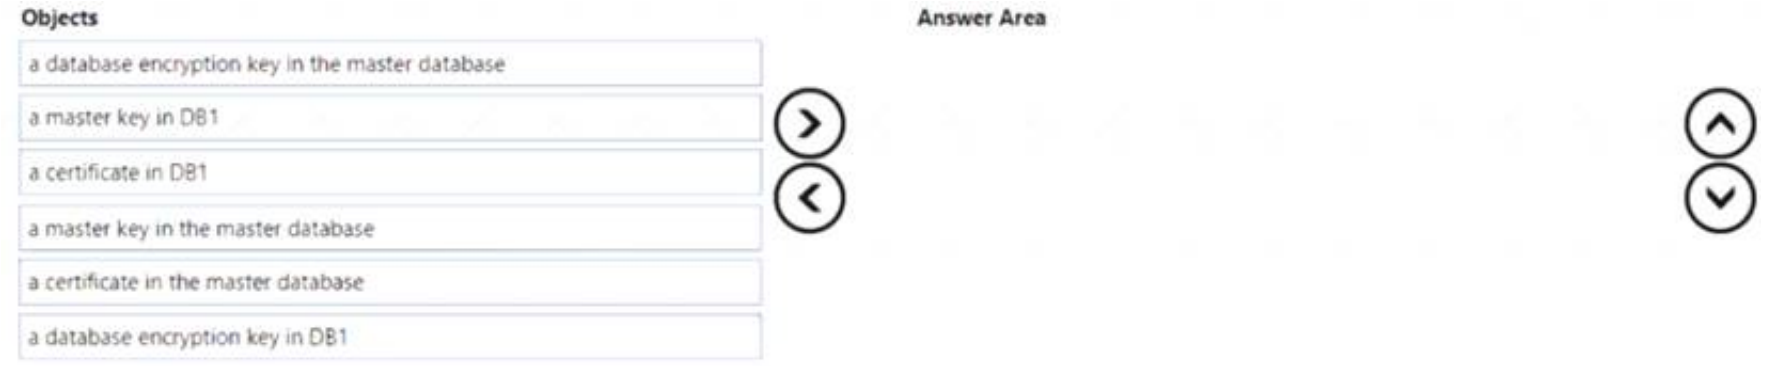

A. Mastered
B. Not Mastered

Answer: A

Explanation:

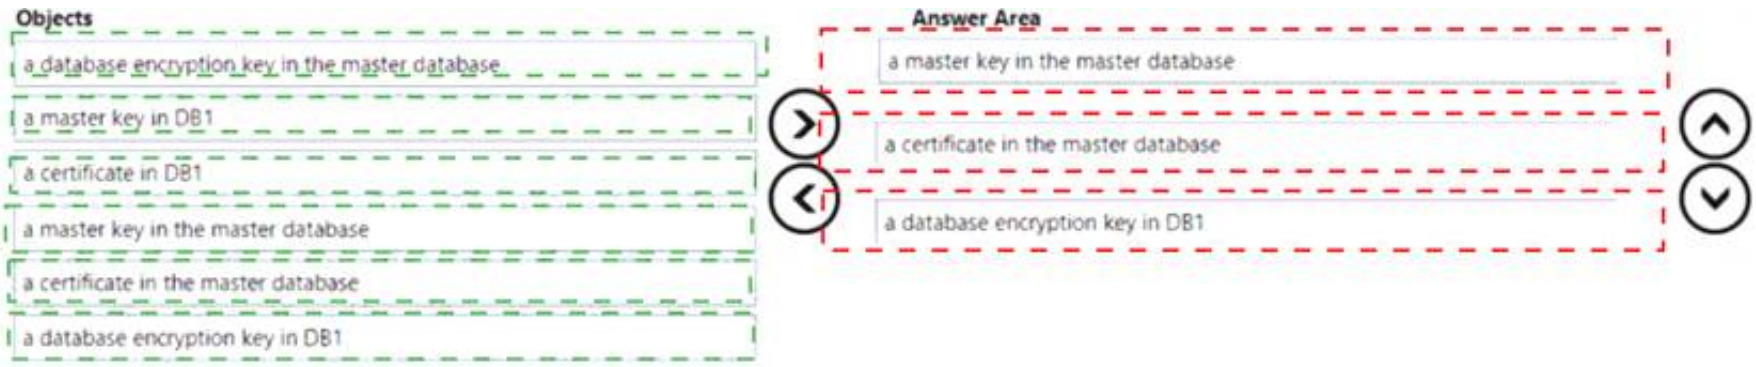

NEW QUESTION 4
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1. You need to identify and repair any physical or logical corruption in DB1. The solution must meet the
following requirements:
• Minimize how long it takes to complete the procedure.
• Minimize data loss.
How should you complete the command? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

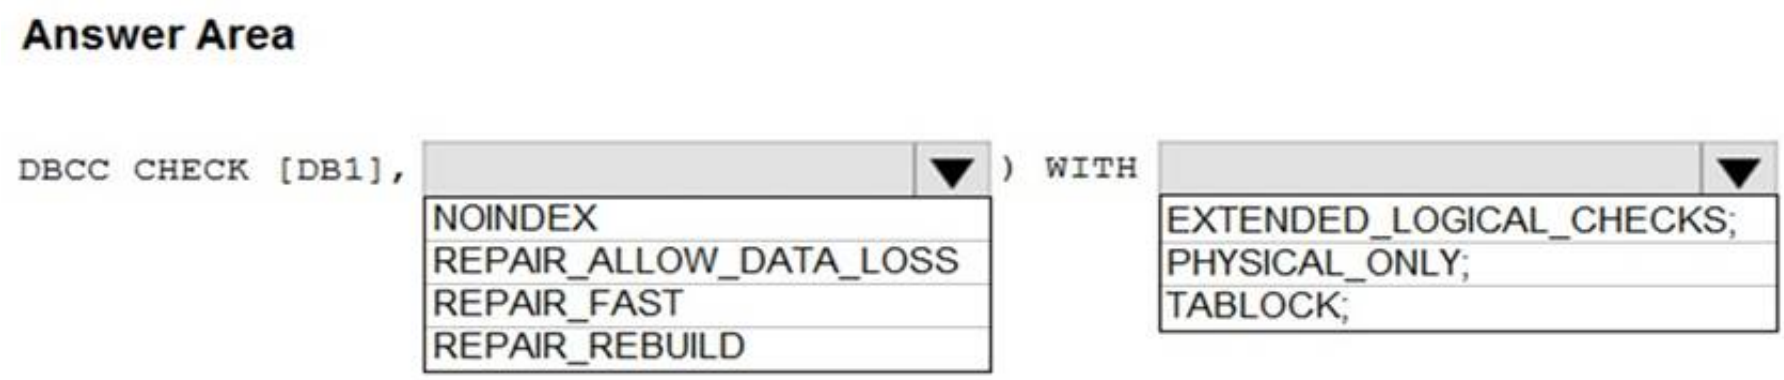

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 5
- (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 6
- (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 7
- (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 8
- (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 9
- (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 querysys.dm_exec_requestsand discover that the wait type isPAGELATCH_UPand thewait_resourceis2:3:905856.
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 10
- (Exam Topic 5)
You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases. You need to update the column and index statistics for the databases.
What should you use?

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

Answer: A

Explanation:

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

NEW QUESTION 10
- (Exam Topic 5)
You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8. You need to deploy a failover cluster instance (FCI) to VM1 and VM2. The solution must eliminate the need for the following:
• A distributed network name (DNN)
• A load balancer What should you do?

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

Answer: D

NEW QUESTION 11
- (Exam Topic 5)
You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:
- Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: C

Explanation:
A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.
The service could be an Azure service such as:
- Azure Storage
- Azure Cosmos DB
- Azure SQL Database
- Your own service using a Private Link Service. Reference:
https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview

NEW QUESTION 15
- (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.

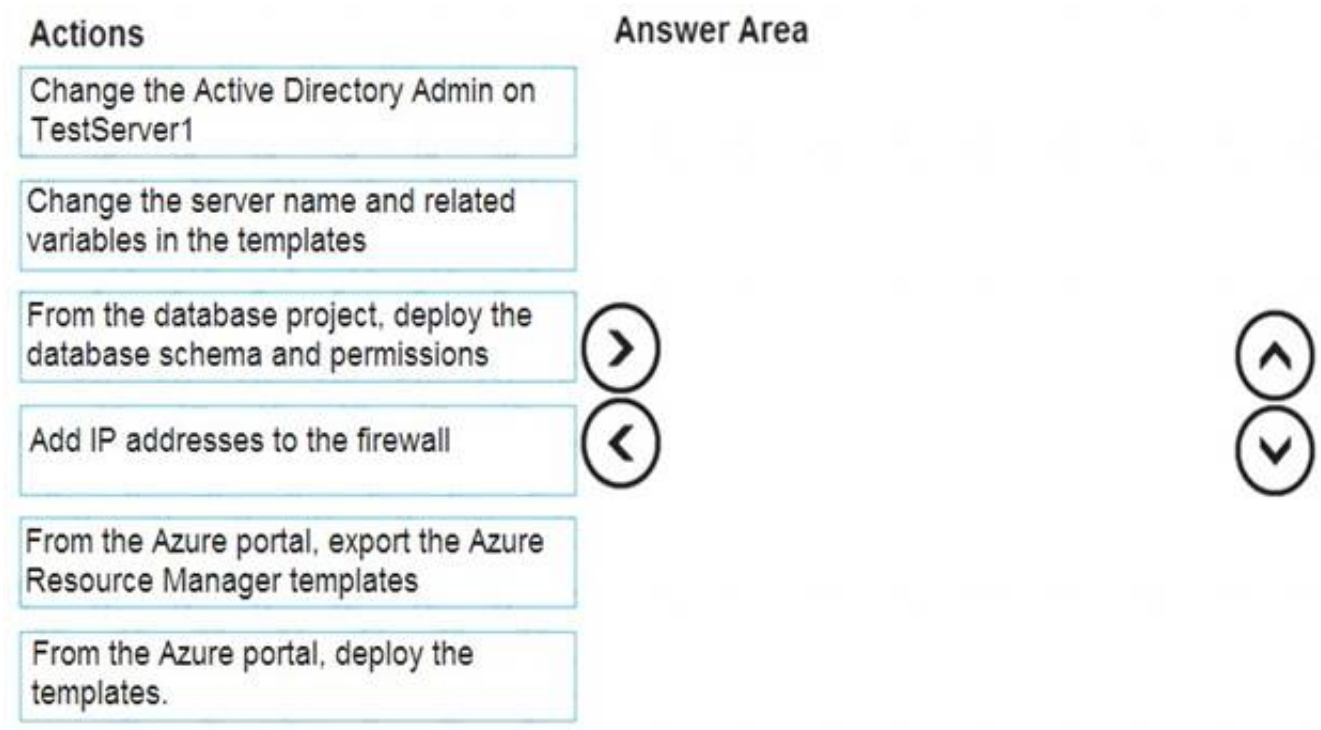

A. Mastered
B. Not Mastered

Answer: A

Explanation:

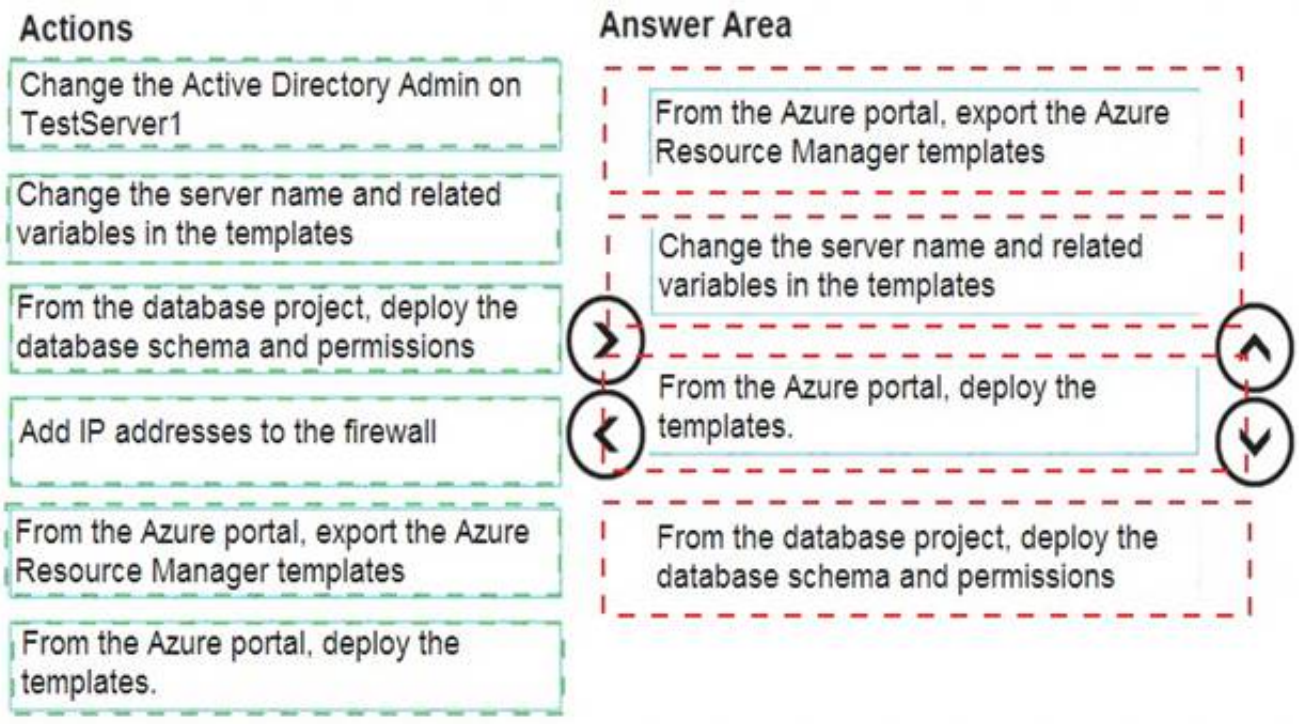

NEW QUESTION 16
- (Exam Topic 5)
You have 10 Azure virtual machines that have SQL Server installed.
You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups.
What should you include in the backup strategy?

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

Answer: B

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

NEW QUESTION 19
- (Exam Topic 5)
You have a Microsoft SQL Server 2017 server that hosts five databases. You Plan to migrate the databases to Azure.
You need to recommend a solution that meets the following requirements:
- Automatically scales compute based on the workload demand
- Provides per-second billing

What should you include in the recommendation? To answer, select the appropriate options in the answer area.

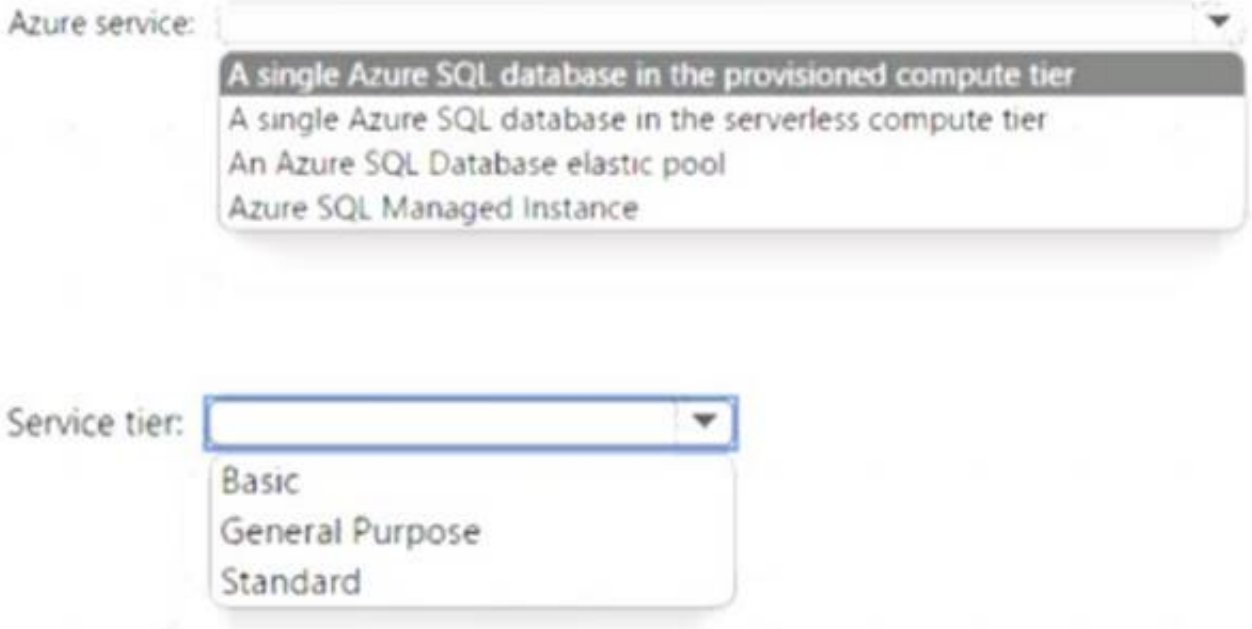

A. Mastered
B. Not Mastered

Answer: A

Explanation:

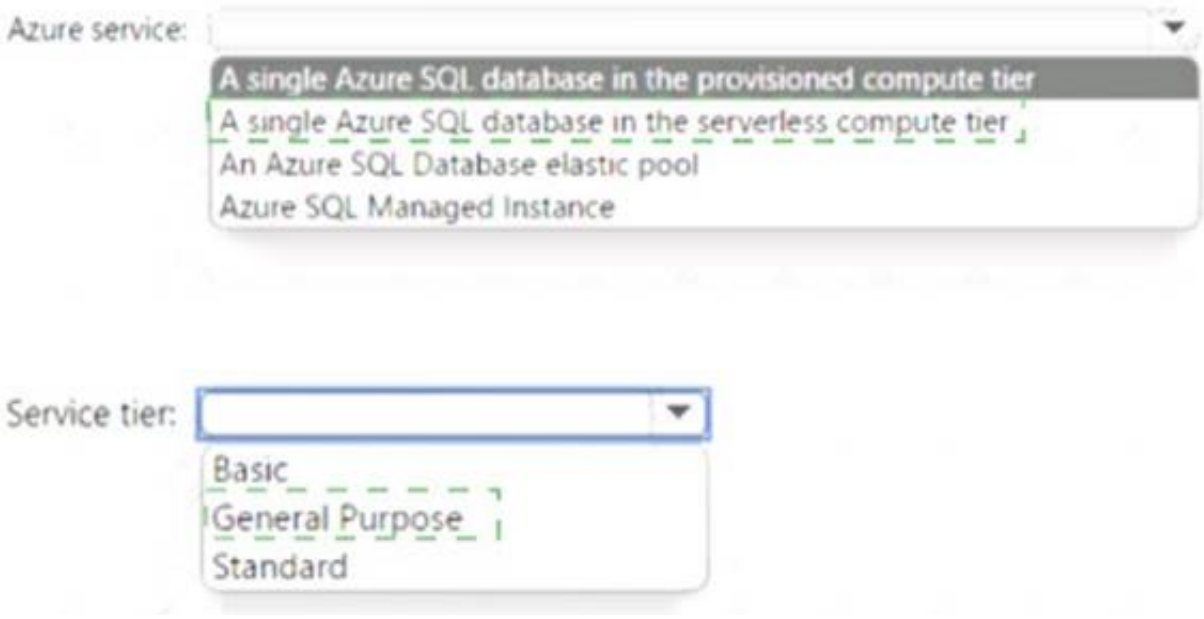

NEW QUESTION 20
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct 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 an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.
Does this meet the goal?

A. Yes
B. No

Answer: A

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

NEW QUESTION 23
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1. You need to modify the MAXDOP settings for db1.
What should you do?

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

Answer: D

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

NEW QUESTION 25
- (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 29
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.
You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed. You need to ensure that webapp1 can identity the changes to Table'. The solution must meet the following
requirements:
• Minimize compute times.
• Minimize storage.
Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

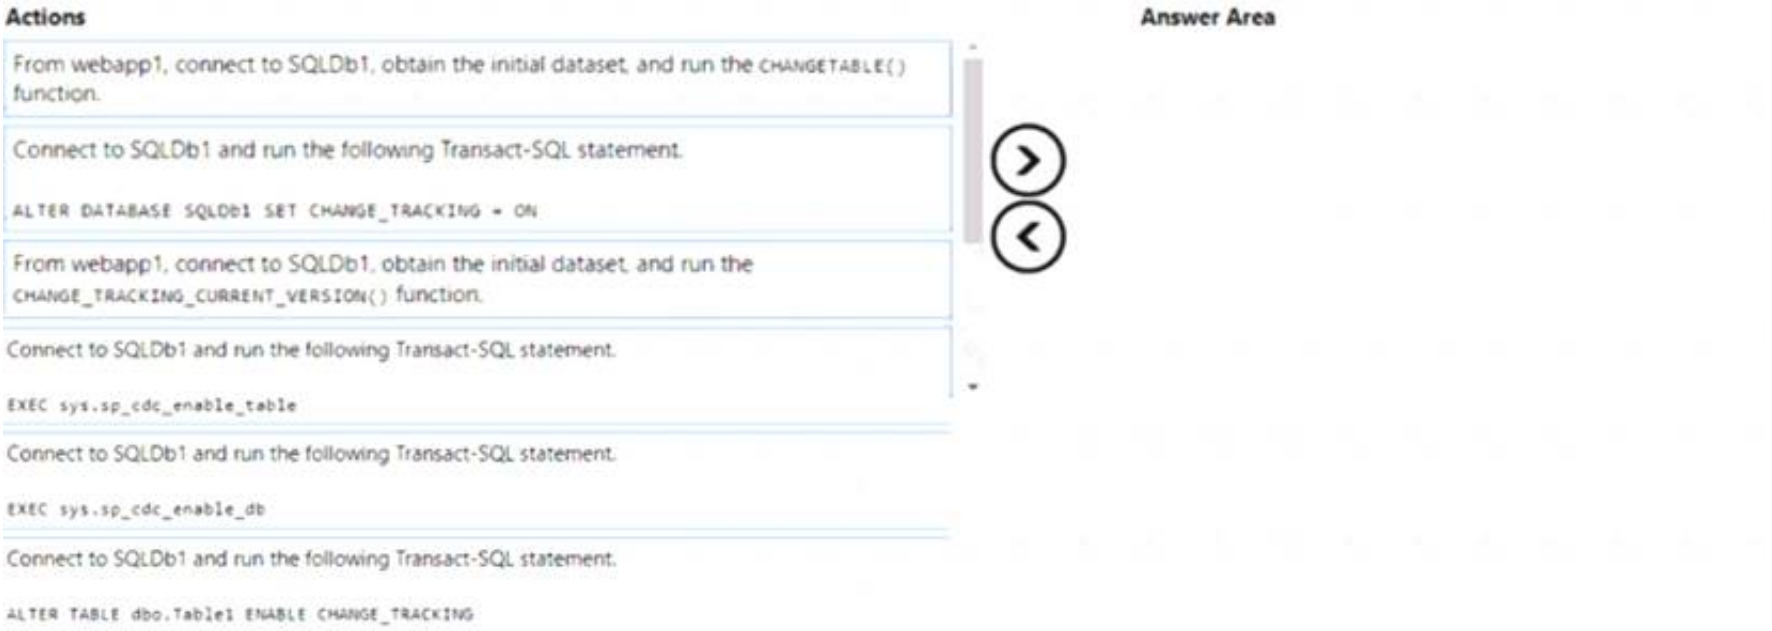

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 30
- (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 33
- (Exam Topic 5)
You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.
You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:
The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
Administrative effort must be minimized. What should you use to host the databases?

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

Answer: D

NEW QUESTION 38
- (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 40
- (Exam Topic 5)
You have SQL Server on an Azure virtual machine that contains a database named DB1. You have an application that queries DB1 to generate a sales report.
You need to see the parameter values from the last time the query was executed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: AC

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

NEW QUESTION 42
- (Exam Topic 5)
You have an Azure subscription.
You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.
You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs. The solution must minimize administrative effort.
Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

MSX: ▼
- SQL database
- SQL managed instances
- SQL virtual machines

TSX: ▼
- SQL database
- SQL managed instances
- SQL virtual machines

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

MSX: ▼
- SQL database
- SQL managed instances
- SQL virtual machines

TSX: ▼
- SQL database
- SQL managed instances
- SQL virtual machines

NEW QUESTION 43
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy a logical SQL server by using PowerShell. The solution must ensure that the logical SQL server can create Azure AD users and provide Transparent Data Encryption (TDE} with a customer-managed key.
How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

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

- assignidentity
-federatedclientID
-keyid

- assignidentity
-federatedclientID
-keyid

A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 47
- (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.

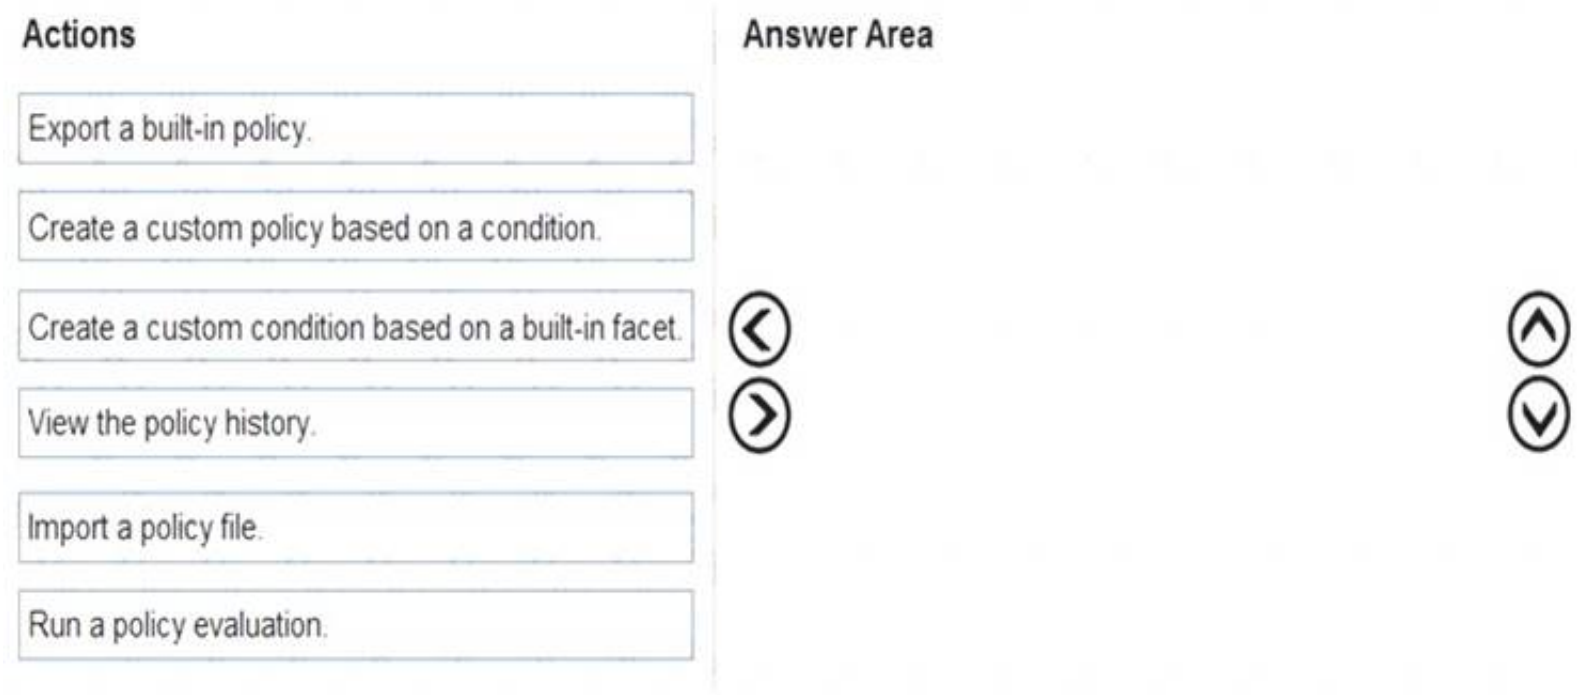

A. Mastered
B. Not Mastered

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 50
- (Exam Topic 5)
You have an Azure subscription.
Vou need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:
• Minimize how long it takes to fail over.
• Maintain existing connections to the primary replica during a failover. What should you do?

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

Answer: C

NEW QUESTION 52
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 database named SQL1 that uses merge replication. You need to migrate SQL1 to Azure. Which service should you use?

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

Answer: C

NEW QUESTION 57
- (Exam Topic 5)
You have an Azure SQL managed instance named MI1.
You need to implement automatic tuning for the databases of MI1. What should you do?

A. Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property.
B. Use Transact-SQL to enable the force_last_good_plan option.
C. From the Azure portal, configure automatic tuning.

Answer: B

NEW QUESTION 60
- (Exam Topic 5)
You have SQL Server 2019 or an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.
Vou scale up the virtual machine to 8 vCPUs and 64 GB of memory.
You need to reduce tempdb contention without regatively affecting server performance. What is the number of secondary data files that you should configure for tempdb?

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

Answer: C

NEW QUESTION 63
- (Exam Topic 5)
You have an Azure SQL database.
You need to implement a disaster recovery solution that meets the following requirements:
• Minimizes how long it takes to recover the database if a datacenter fails
• Minimizes administrative effort
What should you include in the solution?

A. Azure Backup
B. active geo-replication
C. Azure Site Recovery
D. auto-failover groups

Answer: D

NEW QUESTION 67
- (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: C

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

NEW QUESTION 72
- (Exam Topic 5)
You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:
- The web app must be hosted on an Azure virtual network.
- The Azure SQL database must be assigned a private IP address.
- The Azure SQL database must allow connections only from the virtual network. You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

A. Azure Private Link
B. a network security group (NSG)
C. a database-level firewall
D. a server-level firewall

Answer: A

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

NEW QUESTION 73
- (Exam Topic 5)
You have SQL Server on Azure virtual machines in an availability group. You have a database named DB1 that isNOTin 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 75
- (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 80
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.
You plan to migrate DB1 to an Azure SQL Database managed instance.
What should you use to minimize downtime and data loss during the migration?

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

Answer: D

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

NEW QUESTION 85
- (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 90
- (Exam Topic 5)
You configure backup for an Azure SQL database as shown in the following exhibit.

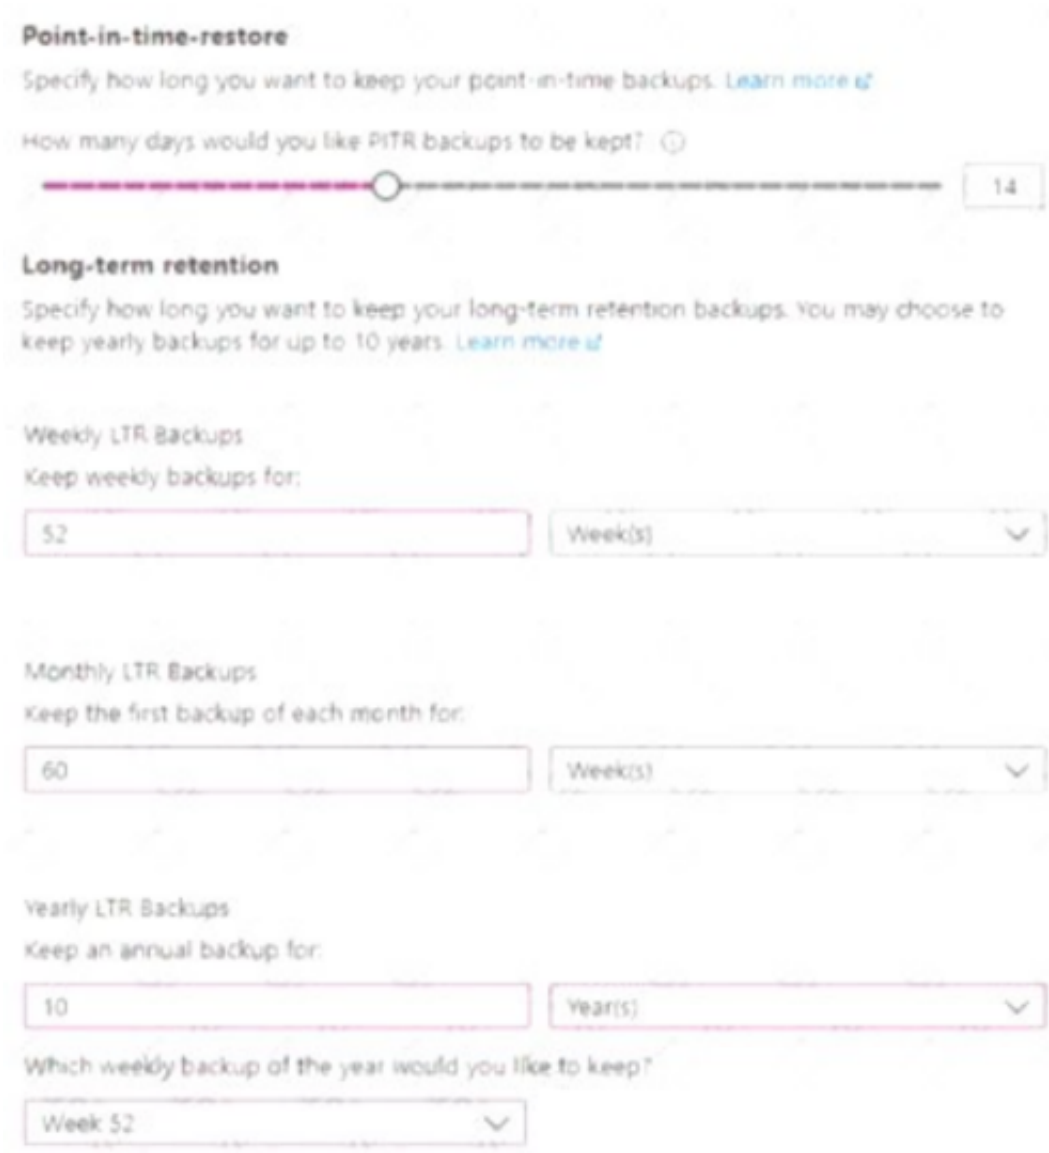

Use the drop-down menus to select the answer choice the completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

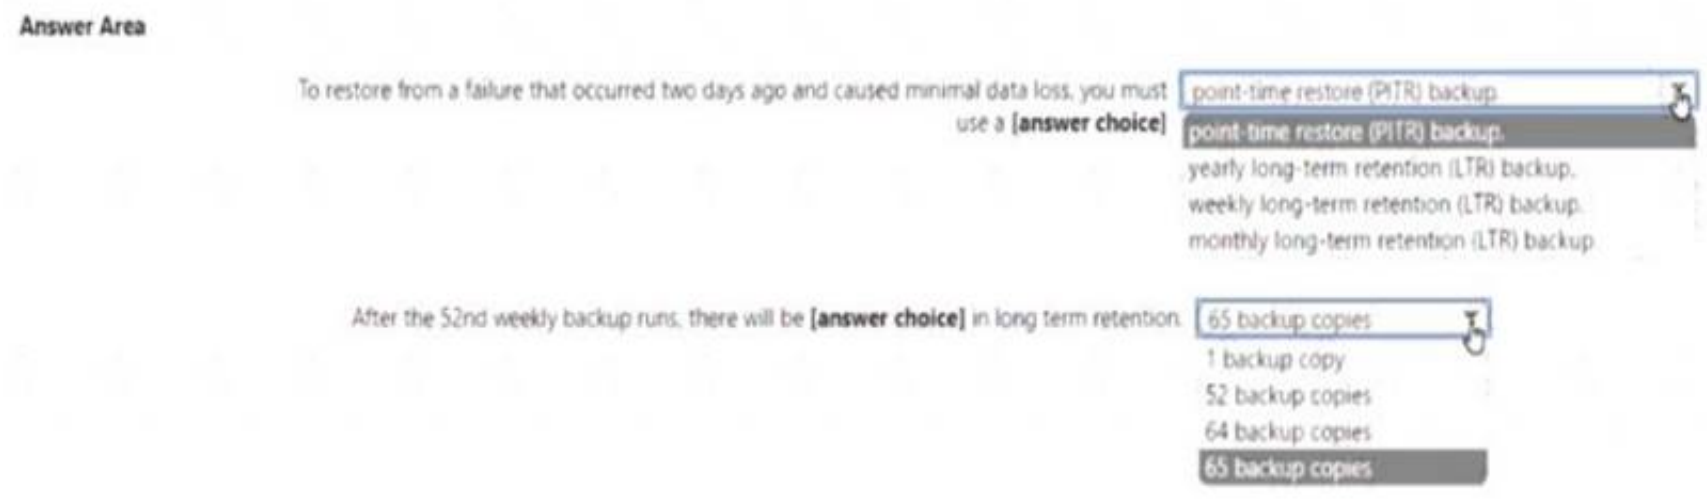

A. Mastered
B. Not Mastered

Answer: A

Explanation:

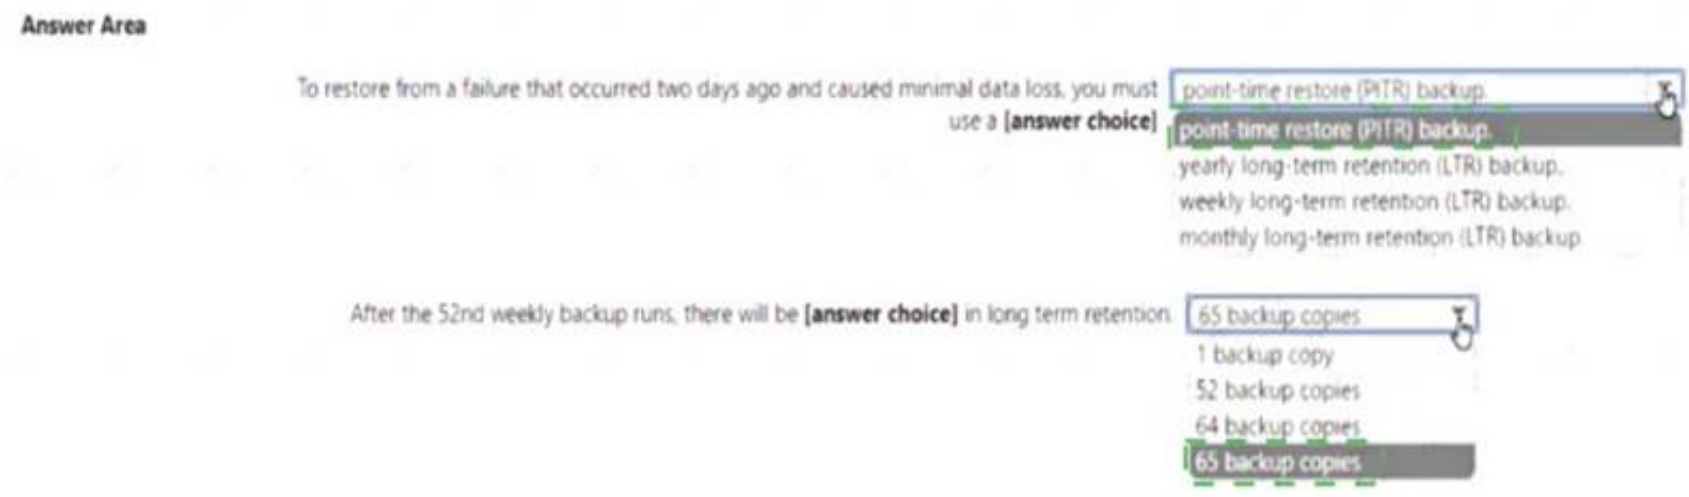

NEW QUESTION 95
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.
You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1.
The on-premises network connects 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 99
- (Exam Topic 5)
You have an Azure Synapse Analytics Apache Spark pool named Pool1.
You plan to load JSON files from an Azure Data Lake Storage Gen2 container into the tables in Pool1. The structure and data types vary by file.
You need to load the files into the tables. The solution must maintain the source data types. What should you do?

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

Answer: B

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

NEW QUESTION 103
- (Exam Topic 5)
You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.
ADF1 contains the following pipelines:

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

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

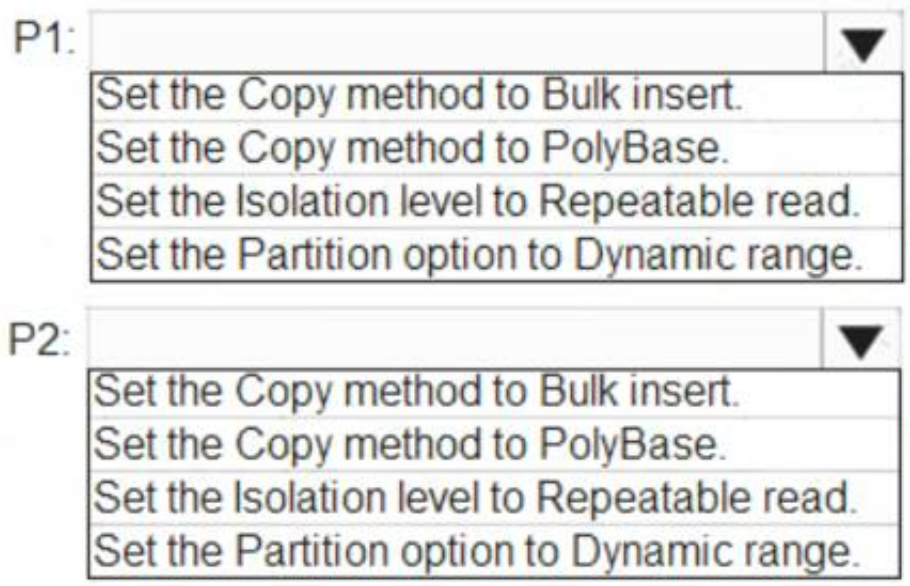

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 105

- (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.

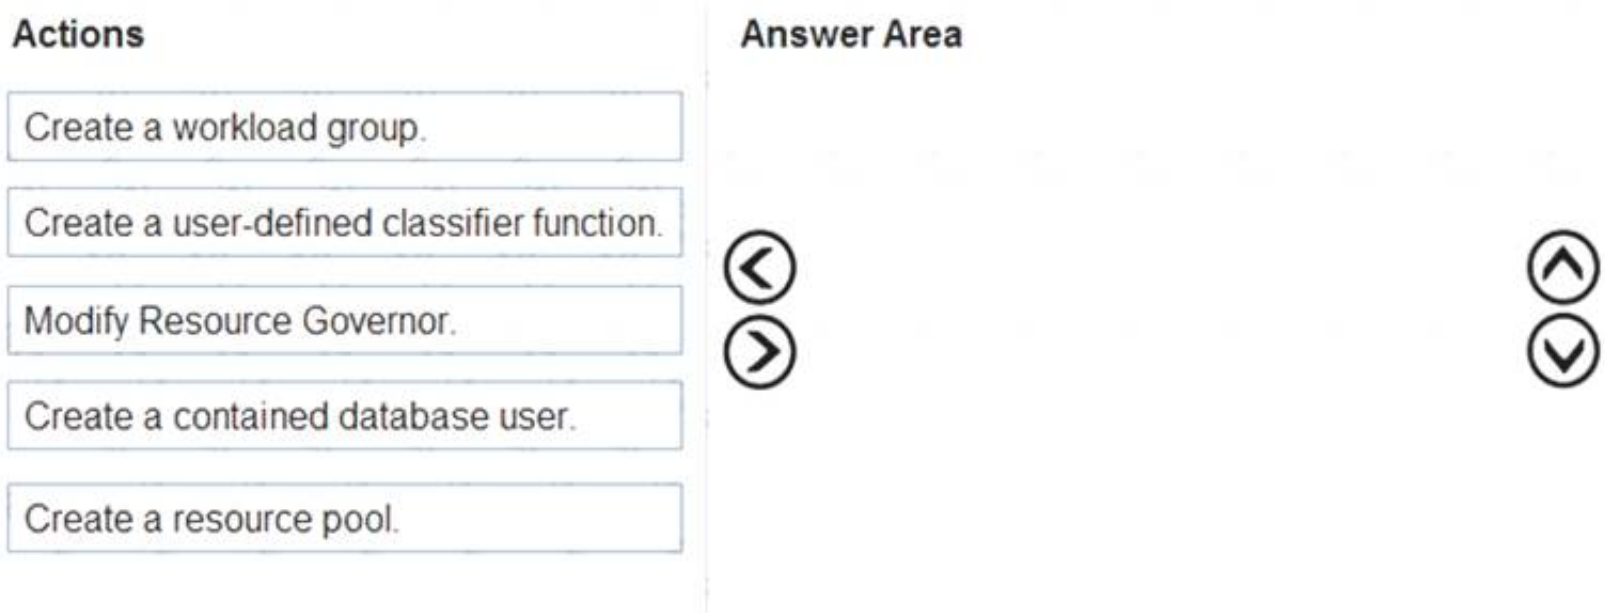

A. Mastered
B. Not Mastered

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 108
- (Exam Topic 5)
You have an Azure SQL database.
You run the following PowerShell script.

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

Connect-AzAccount

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

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

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

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

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 111
- (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.

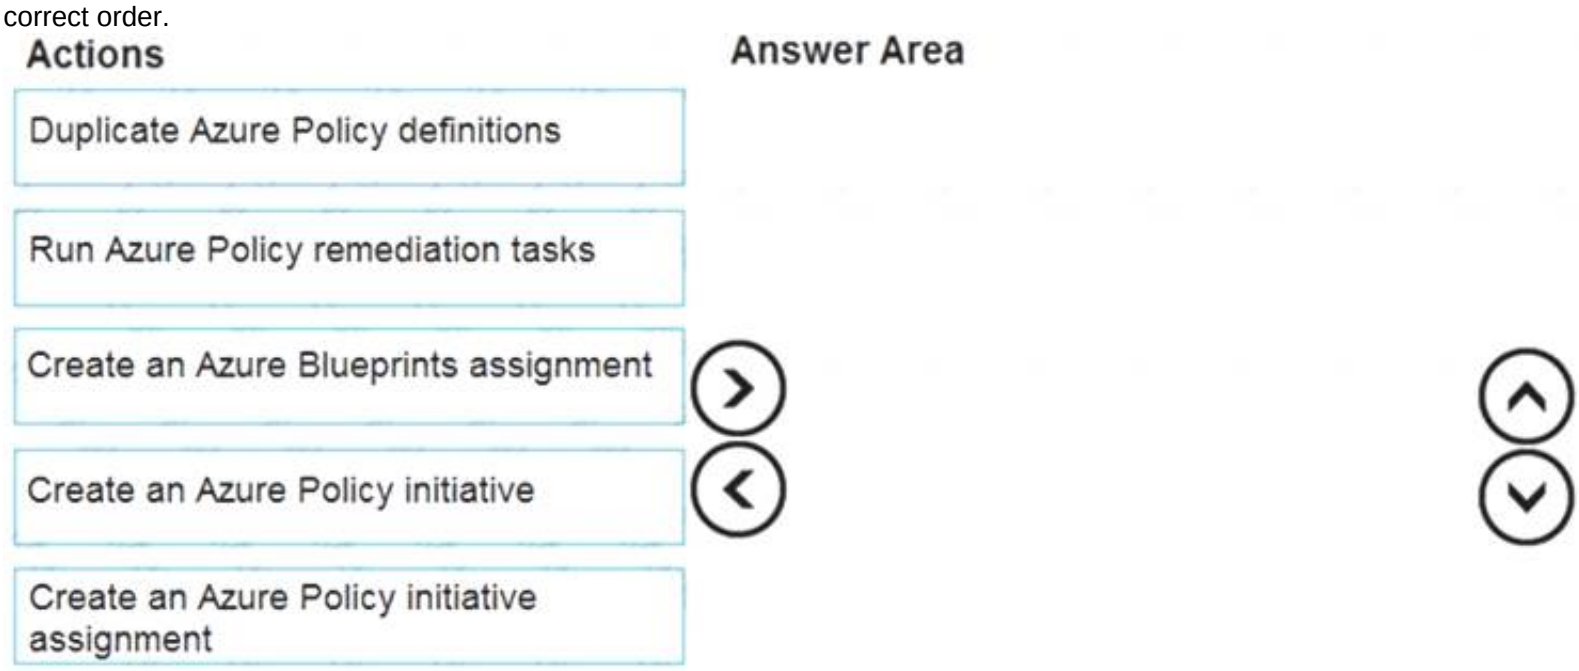

A. Mastered
B. Not Mastered

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 114
- (Exam Topic 5)
You have an Azure SQL database named DB3.
You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click theExhibittab.)

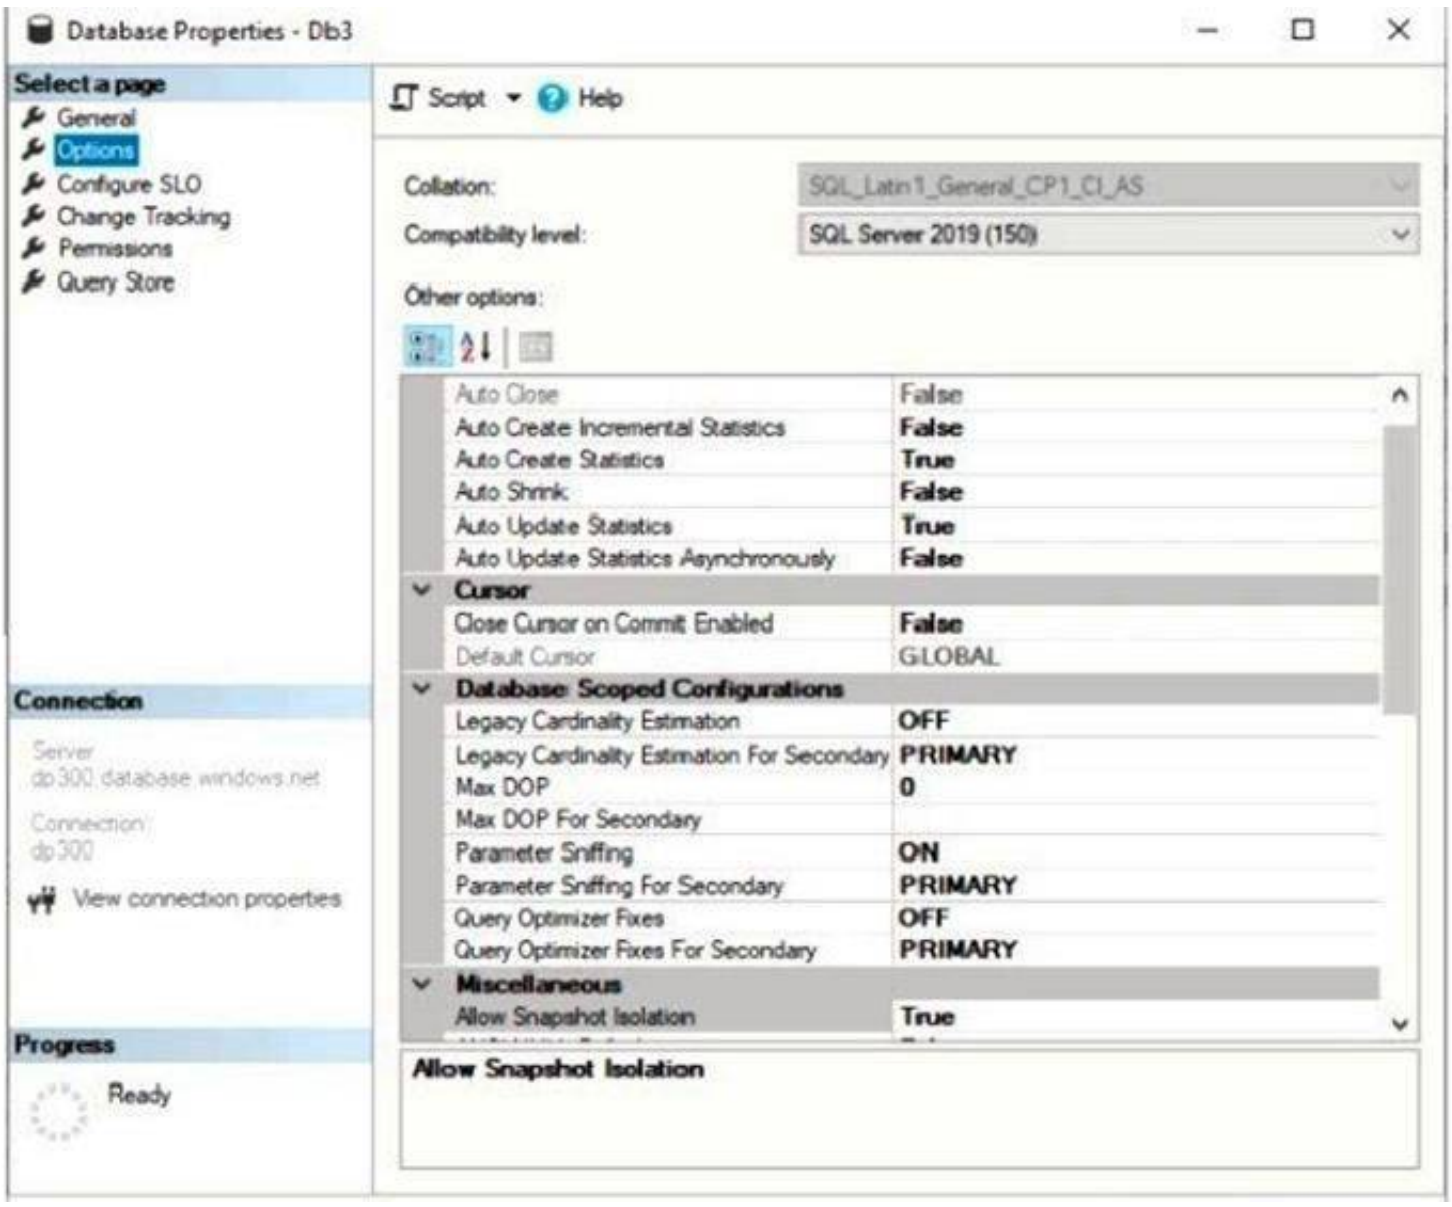

Which Transact-SQL command should you run?

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

Answer: C

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

NEW QUESTION 118
- (Exam Topic 5)
You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:
• Support real-time data replication to a different geographic region.
• Use Azure as a disaster recovery target.
• Minimize costs and administrative effort.
What should you include in the recommendation?

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

Answer: D

NEW QUESTION 119
- (Exam Topic 5)
You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.
You plan to migrate DB1 to an Azure SQL Database managed instance.
What should you use to minimize downtime and data loss during the migration?

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

Answer: D

NEW QUESTION 121
- (Exam Topic 5)

You are creating a managed data warehouse solution on Microsoft Azure.
You must use PolyBase to retrieve data from Azure Blob storage that resides in parquet format and load the data into a large table called FactSalesOrderDetails.
You need to configure Azure Synapse Analytics to receive the data.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

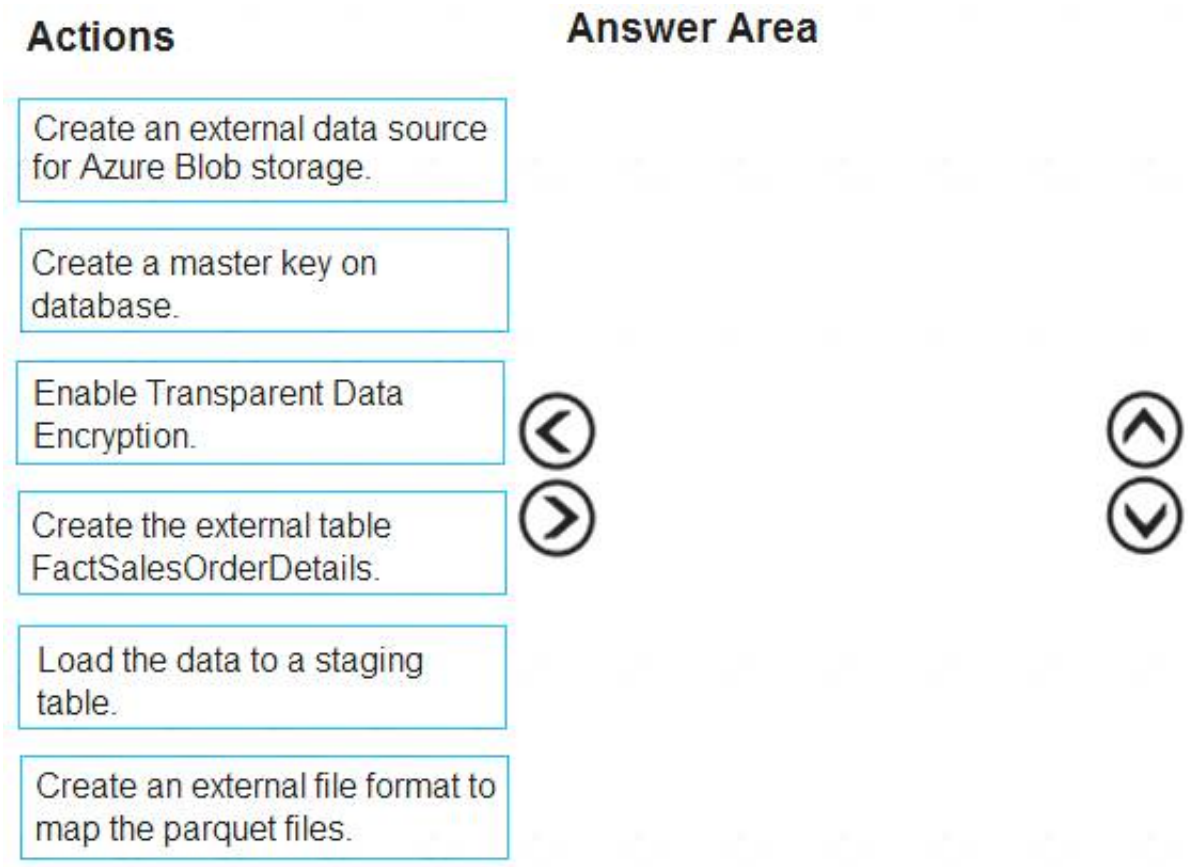

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, chat or text message Description automatically generated
To query the data in your Hadoop data source, you must define an external table to use in Transact-SQL queries. The following steps describe how to configure the external table.
Step 1: Create a master key on database.
* 1. Create a master key on the database. The master key is required to encrypt the credential secret. (Create a database scoped credential for Azure blob storage.)
Step 2: Create an external data source for Azure Blob storage.
* 2. Create an external data source with CREATE EXTERNAL DATA SOURCE.. Step 3: Create an external file format to map the parquet files.
* 3. Create an external file format with CREATE EXTERNAL FILE FORMAT. Step 4. Create an external table FactSalesOrderDetails
* 4. Create an external table pointing to data stored in Azure storage with CREATE EXTERNAL TABLE. Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-configure-azure-blob-storage

NEW QUESTION 122
- (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 127
- (Exam Topic 5)
You have an Azure SQL database named DBI that contains a nonclustered index named index1. End users report slow queries when they use index1.
You need to identify the operations that are being performed on the index. Which dynamic management view should you use?

A. sys.dm_exec_query_plan_stats

B. Sys.dm_db_index_physical_stats

C. Sys.dm_db_index_operational_stats

D. Sys.dm_db_index_useage_stats

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

Answer: D

NEW QUESTION 129
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named MI1 and an Azure Storage account named storage1.
You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:
* Use block blob storage.
* Maximize security.

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

Answer: B

NEW QUESTION 133
- (Exam Topic 5)
You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.
You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.
You need to recommend a database service for the deployment. The solution must minimize administrative effort.
What should you include in the recommendation?

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

Answer: B

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

NEW QUESTION 137
- (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 142
- (Exam Topic 3)
Which windowing function should you use to perform the streaming aggregation of the sales data?
A. Sliding
B. Hopping
C. Session
D. Tumbling

Answer: D

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

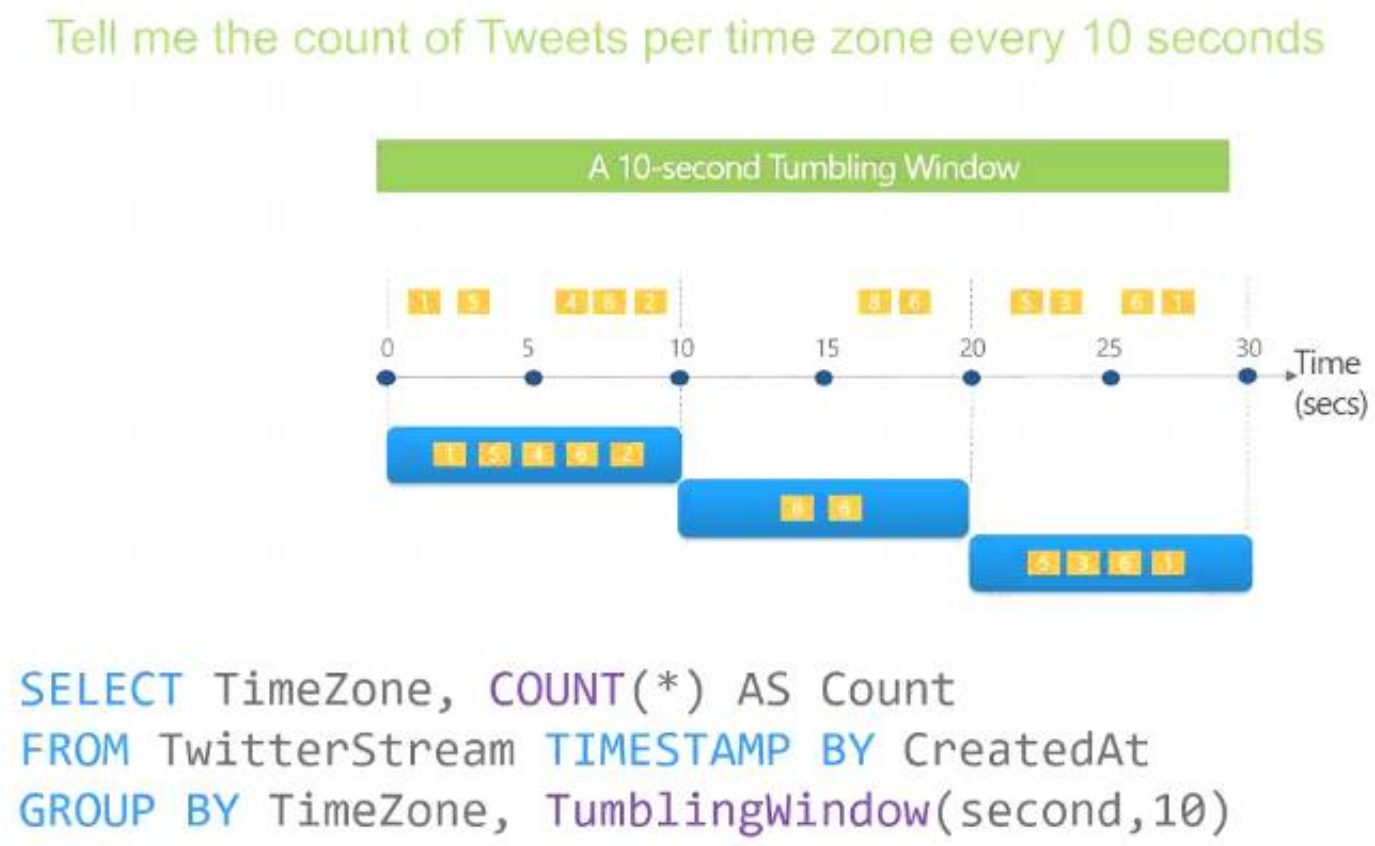

Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 145
- (Exam Topic 1)
You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements. How should you collect and stream metrics? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

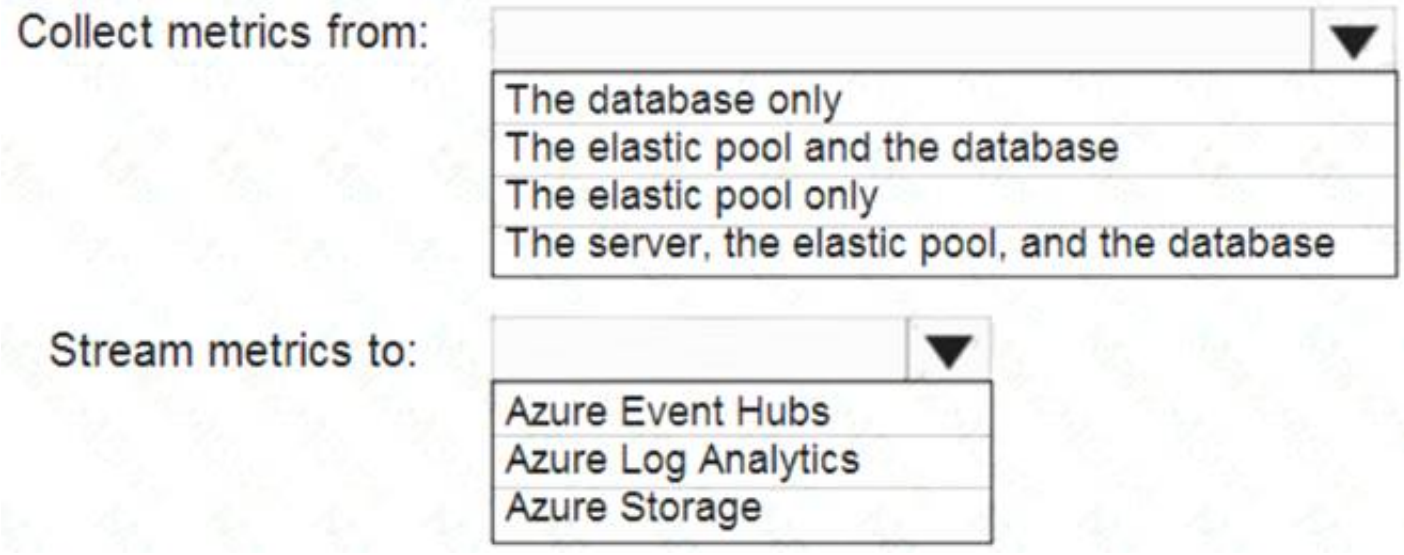

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: The server, the elastic pool, and the database
Senario:
SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.
Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.
Box 2: Azure Event hubs
Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform. Event hubs are able to handle custom metrics.

NEW QUESTION 146
- (Exam Topic 1)
You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.
What should you include in the plan?

A. Configure the Deleted databases settings for ResearchSrvOL
B. Deploy and configure an Azure Backup server.
C. Configure the Advanced Data Security settings for ResearchDBL
D. Configure the Manage Backups settings for ResearchSrvOL

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/long-term-backup-retention-configure

NEW QUESTION 151
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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