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

```
EXECUTE AS         ▼   = 'user1@contoso.com'
                CALLER
                LOGIN
                OWNER
                USER

GO


SELECT  SUSER_SNAME()

        ▼
  REVERT
  REVOKE
  ROLLBACK

GO
```

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 2
- (Exam Topic 5)
You have an Azure subscription that uses a domain named contoso.com.
You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.
DBServer1 has a high volume of data changes and low latency requirements for data writes.
You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.
What should you do?

A. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous.
B. Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous.
C. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous.
D. Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous.

Answer: C

Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on

NEW QUESTION 3
- (Exam Topic 5)
You have an Azure subscription.
You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.
You need to configure the disks on the virtual machine. The solution must meet the following requirements:
• Minimize latency for transaction logs.
• Minimize the impact on IO Of the virtual machine.
Which type of disk should you use for each workload? To answer. drag the appropriate disk types to the correct workloads. Each disk type may be used once, 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.

Disk Types	Answer Area
Local	TempDB: Disk Type
Premium SSD	Transaction logs: Disk Type
Standard HDD	
Standard SSD	
Ultra Disk	

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 4
- (Exam Topic 5)
You have an on-premises Microsoft SQL Server 2019 instance that hosts a database named DB1.
You plan to perform an online migration of DB1 to an Azure SQL managed instance by using the Azure Database Migration Service.
You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.
What should you run for the backup and where should you store the backup? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Run:
- A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH CHECKSUM option
- A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:
- A Recovery Services vault
- An Azure Blob storage account
- An SMB file share

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 5
- (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 6
- (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
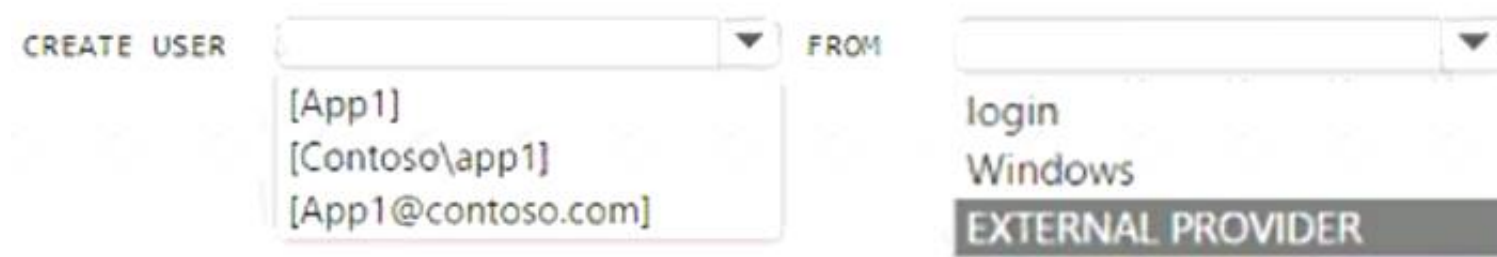D. an Azure Blob storage container

Answer: A

NEW QUESTION 7
- (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 8
- (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: BD

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

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 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 10
- (Exam Topic 5)
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:
• Capture raw event data and store the data in Azure Storage.
• Minimize the performance impact of capturing extended events.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
CREATE   EVENT SESSION session1 ON DATABASE

   ADD EVENT sqlserver.sql_statement_starting

      (

      ACTION (sqlserver.sql_text)

      WHERE statement LIKE 'UPDATE gmTabEmployee%'

      )

ADD TARGET

   package0.    event_file            ▼
                event_file
                event_stream
                ring_buffer

      )

   SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'

      )

   WITH

      (MAX_MEMORY = 10 MB,

      EVENT_RETENTION_MODE=   ALLOW_MULTIPLE_EVENT_LOSS    ▼
                             ALLOW_MULTIPLE_EVENT_LOSS
   MAX_DISPATCH_LATENCY = 3 SE( ALLOW_SINGLE_EVENT_LOSS
                             NO_EVENT_LOSS
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

```
CREATE  EVENT SESSION session1 ON DATABASE

  ADD EVENT sqlserver.sql_statement_starting

    (

    ACTION (sqlserver.sql_text)

    WHERE statement LIKE 'UPDATE gmTabEmployee%'

    )

  ADD TARGET

    package0.    event_file            ▼
                 event_file
                 event_stream
                 ring_buffer

    )

    SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'

    )

  WITH

    (MAX_MEMORY = 10 MB,

    EVENT_RETENTION_MODE=    ALLOW_MULTIPLE_EVENT_LOSS    ▼
                            ALLOW_MULTIPLE_EVENT_LOSS
                            ALLOW_SINGLE_EVENT_LOSS
    MAX_DISPATCH_LATENCY = 3 SE    NO_EVENT_LOSS
```

NEW QUESTION 10
- (Exam Topic 5)
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.
You need to ensure that User1 can perform the following tasks on DB1:
• Create jobs.
• View all jobs.
• Modify, delete, and disable the jobs the user created. The solution must use the principle of least privilege.
Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Location: ▼
             DB1
             master
             msdb

Built-in role: ▼
                  SQLAgentOperatorRole
                  SQLAgentReaderRole
                  SQLAgentUserRole

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 11
- (Exam Topic 5)
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies that several tables are missing indexes. You need to ensure that indexes are created for the tables.
What should you do?

A. RuntheDBCC SQLPERF command.

B. Run thedbcc dbreindexcommand.
C. Modify the automatic tuning settings for db1.
D. Modify the Query Store settings for db1.

Answer: C

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

NEW QUESTION 12
- (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 13
- (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 16
- (Exam Topic 5)
You have an Azure Data Lake Storage Gen2 container.
Data is ingested into the container, and then transformed by a data integration application. The data isNOT modified after that. Users can read files in the container but cannot modify the files.
You need to design a data archiving solution that meets the following requirements:
≫ New data is accessed frequently and must be available as quickly as possible.
≫ Data that is older than five years is accessed infrequently but must be available within one second when requested.
≫ Data that us older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
≫ Costs must be minimized while maintaining the required availability.
How should you manage the data? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Five-year-old data:
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data:
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 19
- (Exam Topic 5)
You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database.
You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance. Which three service tiers support the SQL managed instance? Each correct answer presents a complete
solution. NOTE: Each correct selection is worth one point.

A. General Purpose Standard
B. Business Critical Premium
C. Business Critical Memory Optimized Premium
D. General Purpose Premium
E. Business Critical Standard

Answer: BCD

NEW QUESTION 23
- (Exam Topic 5)
You have an Azure Databricks resource.
You need to log actions that relate to changes in compute for the Databricks resource. Which Databricks services should you log?

A. clusters
B. jobs
C. DBFS
D. SSH
E. workspace

Answer: E

Explanation:
Cloud Provider Infrastructure Logs.
Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace.
Customers, especially in the regulated industries, also need records on activities like:

≫ User access control to cloud data storage

≫ Cloud Identity and Access Management roles

≫ User access to cloud network and compute
Azure Databricks offers three distinct workloads on several VM Instances tailored for your data analytics workflow—the Jobs Compute and Jobs Light Compute workloads make it easy for data engineers to build and execute jobs, and the All-Purpose Compute workload makes it easy for data scientists to explore, visualize, manipulate, and share data and insights interactively.
Reference:
https://databricks.com/blog/2020/03/25/trust-but-verify-with-databricks.html

NEW QUESTION 24
- (Exam Topic 5)
You are building a database in an Azure Synapse Analytics serverless SQL pool. You have data stored in Parquet files in an Azure Data Lake Storage Gen2 container. Records are structured as shown in the following sample.

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

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

NOTE:Each correct selection is worth one point.

```
                          ▼  applications
CREATE EXTERNAL TABLE
CREATE TABLE
CREATE VIEW
WITH (
    LOCATION = 'applications/',
    DATA_SOURCE = applications_ds,
    FILE_FORMAT = applications_file_format
)
AS
SELECT id, [address_housenumber] as addressnumber, [address_line1]
as addressline1
FROM
              ▼   (BULK 'https://contoso1.dfs.core.windows.net/
CROSS APPLY       applications/year=*/*.parquet',
OPENJSON
OPENROWSET
       FORMAT = 'PARQUET') AS [r]
  GO
```

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 25
- (Exam Topic 5)
You have an Azure SQL database named db1.
You need to retrieve the resource usage of db1 from the last week.
How should you complete the statement? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
SELECT *

FROM              _____ ▼
                  sys.dm_db_resource_stats
                  sys.dm_exec_requests
                  sys.dm_user_db_resource_governance
                  sys.resource_stats

WHERE database_name = 'db1' AND

    start_time >   _____ ▼   (day, -7, GETDATE())
                   DATEADD
                   DATEDIFF
                   DATEPART
                   TODATETIMEOFFSET

ORDER BY start_time DESC;
```

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 28
- (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 30
- (Exam Topic 5)
You have an Azure SQL managed instance.
You need to restore a database named DB1 by using Transact-SQL.
Which command should you run? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
RESTORE  [ ▼ ]  DB1  FROM  [                                          ▼ ]
         DATABASE           DISK = N'\\NAS01\$SQLBackups\DB1.bak';
         FILE               TAPE = N'\\.\Tape0'
         LOG                URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated

NEW QUESTION 33
- (Exam Topic 5)
You have a single availability set that contains two SQL Server on Azure Virtual Machines instances.
The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1.
You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:
• Provide an availability SLA
• Minimize costs.
What should you create?

A. a virtual network name (VNN) resource
B. a Basic Azure Load Balancer
C. a distributed network name (DNN) resource
D. an Azure Standard Load Balancer

Answer: C

NEW QUESTION 34

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

ALTER DATABASE [Db1]

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

GO
ALTER DATABASE [Db1]

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

GO

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 39
- (Exam Topic 5)
You have an Azure subscription that contains a SQL Server on Azure Virtual Machines instance named SQLVMI. SQLVMI hosts a database named OBI.
You need to retrieve query plans from the Query Store on DBI. What should you do first?

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

Answer: B

NEW QUESTION 42
- (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 47
- (Exam Topic 5)
You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service.
You have an Azure Key vault named vault1 that contains an encryption kay named key1. You need to encrypt df1 by using key1.
What should you do first?

A. Disable purge protection on vault1.
B. Remove the linked service from df1.
C. Create a self-hosted integration runtime.
D. Disable soft delete on vault1.

Answer: B

Explanation:
A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.
Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key

NEW QUESTION 51
- (Exam Topic 5)
You have an Azure subscription that contains 50 instances of SQL Server on Azure Virtual Machines. The instances host 500 Azure SQL databases. You need to ensure that all the databases have the same configuration. The solution must meet the following requirements:
• Auditing must be enabled.
• Azure Defender must be enabled.
• Public network access must be disabled.
• Administrative effort must be minimized.
Which two resources should you create in the subscription? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. an Azure Policy assignment
B. an Azure Automation account
C. an Azure Policy initiative
D. an Azure Automation runbook
E. an Azure Policy definition

Answer: CE

NEW QUESTION 56
- (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 57
- (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.

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Azure service:

A single Azure SQL database in the provisioned compute tier
A single Azure SQL database in the serverless compute tier
An Azure SQL Database elastic pool
Azure SQL Managed Instance

Service tier:

Basic
General Purpose
Standard

NEW QUESTION 61
- (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 66
- (Exam Topic 5)
You have an Azure SQL Database instance named DatabaseA on a server named Server1.
You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.
You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

On the master database, run `CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;`

On DatabaseA, run `CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';`

On DatabaseA, run `ALTER ROLE db_datareader ADD MEMBER [App1];`

On the master database, run `CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';`

On DatabaseA, run `CREATE USER [App1] FROM LOGIN [App1];`

Answer Area

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 67
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.
Which deployment option and pricing tier should you configure?

A. Azure SQL Database Hyperscale
B. Azure SQL Database managed instance General Purpose
C. Azure SQL Database Premium
D. Azure SQL Database Basic

Answer: C

Explanation:
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

NEW QUESTION 70
- (Exam Topic 5)
You have an Azure subscription that contains an Azure SQL database named SQL1. SQL1 is in an Azure region that does not support availability zones.
You need to ensure that you have a secondary replica of SQLI in the same region. What should you use?

A. log shipping
B. auto-failover groups
C. active geo-replication
D. Microsoft SQL Server failover clusters

Answer: C

NEW QUESTION 73
- (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"}   [ - assignidentity ▼ ]   [ - assignidentity ▼ ]
                                                                  [ -federatedclientID ]     [ -federatedclientID ]
                                                                  [ -keyid           ]       [ -keyid           ]

    "https://db1.vault.azure.net/keys/dbkey/012342345123456789012345618 23942"
```

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 75
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains a private certificate named Sales. The private key for Sales is encrypted with a password. You need to change the password for the private key. Which
Transact-SQL statement should you run?
A)

```
ALTER CERTIFICATE Sales
    WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%',
    ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```

B)

```
ALTER CERTIFICATE Sales
    WITH PRIVATE KEY (ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```

C)

```
ALTER CERTIFICATE Sales    WITH PRIVATE KEY (FILE = 'D:\importkeys\SalesNew,    DECRYPTION BY PASSWORD = ' Mb^6BK&*w%');
```

D)

```
ALTER CERTIFICATE Sales    WITH PRIVATE KEY (DECRYPTION BY PASSWORD = ' EWYx9Xk+$#');
```

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

Answer: C

NEW QUESTION 78
- (Exam Topic 5)
You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.
You need to recommend an authentication solution for Pool1. The solution must support multi-factor authentication (MFA) and database-level authentication.
Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

MFA: [dropdown]
| Azure AD authentication |
| Microsoft SQL Server authentication |
| Passwordless authentication |
| Windows authentication |

Database-level authentication: [dropdown]
| Application roles |
| Contained database users |
| Database roles |
| Microsoft SQL Server logins |

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 80
- (Exam Topic 5)
You have an Azure SQL managed instance.
You need to enable SQL Agent Job email notifications. What should you do?

A. Use the Agent XPs option.
B. Enable the SQL Server Agent.
C. Run the sp_configure command.
D. Run the sp set_agent properties command.

Answer: C

NEW QUESTION 82
- (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 all traffic to the public endpoint of SqlSrv1 is blocked.
➢ Minimize the possibility of VM1 exfiltrating data stored in SqlDb1. What should you create on VNet1?

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

Answer: C

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

NEW QUESTION 85
- (Exam Topic 5)
You have an Azure SQL database.
You need to identify whether a delayed query execution is associated to a RESOURCE wait.
How should you complete the Transact –SQL statement? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

NEW QUESTION 88
- (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 90
- (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 run theRemove-AzSqlDatabasePowerShell cmdlet for Database1 on Server2. You run theRestore-AzSqlDatabasePowerShell cmdlet for Database1 on Server2.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.
Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

NEW QUESTION 95
- (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 96
- (Exam Topic 5)
You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.
You need to protect the contents of Column1 at rest, in transit, and in use.
How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Encryption key:
- Column encryption key
- Database encryption key
- Service master key

Encryption type:
- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Column encryption Key
Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine

NEW QUESTION 98
- (Exam Topic 5)
You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.
You need to implement alerts by using Azure Monitor. The solution must meet the following requirements: ≫ Minimize costs.

≫ Aggregate Intelligent Insights telemetry from each database. What should you do?

A. From the Diagnostic settings of each database, select Send to Log Analytics.
B. From the Diagnostic settings of each database, select Stream to an event hub.
C. From the Diagnostic settings of SQLMI1. select Send to Log Analytics.
D. From the Diagnostic settings of SQLMI1. select Stream to an event hub.

Answer: A

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo

NEW QUESTION 103

- (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 108
- (Exam Topic 5)
You have an Azure SQL Database elastic pool that contains 10 databases. You receive the following alert.
Msg 1132, Level 16, State 1, Line 1
The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.
You need to resolve the alert. The solution must minimize administrative effort.
Which three actions can you perform? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

Answer: BCD

NEW QUESTION 113
- (Exam Topic 5)
You have a SQL Server on Azure Virtual Machines instance named VM1 . You run the following query.

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

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

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 114
- (Exam Topic 5)
You configure backups for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore

Specify how long you want to keep your point-in-time backups. Learn more ☑

How many days would you like PITR backups to be kept? ⓘ

| 14 |

Long-term retention

Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. Learn more ☑

Weekly LTR Backups
Keep weekly backups for:

| 52 | | Week(s) ∨ |

Monthly LTR Backups
Keep the first backup of each month for:

| 60 | | Week(s) ∨ |

Yearly LTR Backups
Keep an annual backup for:

| 10 | | Year(s) ∨ |

Which weekly backup of the year would you like to keep?

| Week 52 ∨ |

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

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

| ▼ |
| point-time restore (PITR) backup. |
| yearly long-term retention (LTR) backup. |
| weekly long-term retention (LTR) backup. |
| monthly long-term retention (LTR) backup. |

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

| ▼ |
| 1 backup copy |
| 52 backup copies |
| 64 backup copies |
| 65 backup copies |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice]

| ▼ |
| point-time restore (PITR) backup. |
| yearly long-term retention (LTR) backup. |
| weekly long-term retention (LTR) backup. |
| monthly long-term retention (LTR) backup. |

After the 52nd weekly backup runs, there will be [answer choice] in long term retention.

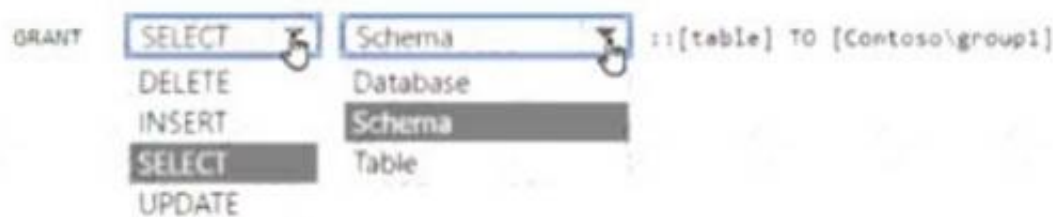| ▼ |
| 1 backup copy |
| 52 backup copies |
| 64 backup copies |
| 65 backup copies |

NEW QUESTION 116
- (Exam Topic 5)
You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created m 06l The solution must use the principle of least privilege Howshould you complete the Transact-SQL statement' To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 120
- (Exam Topic 5)
You manage an enterprise data warehouse in Azure Synapse Analytics.
Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.
You need to monitor resource utilization to determine the source of the performance issues. Which metric should you monitor?

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

Answer: A

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

NEW QUESTION 121
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named
container1.
You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.
You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.
Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead use a serverless SQL pool to create an external table with the extra column. Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

NEW QUESTION 123
- (Exam Topic 5)
You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

Name	Type
Column1	Ntext
Column2	Geometry
Column3	Image
Column4	Varchar
Column5	Datetime2

You plan to enable Always Encrypted for Table1.
Which two columns support encryption? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point

A. Column1
B. Column2
C. Column3
D. Column4
E. Column5

Answer: AD

NEW QUESTION 128
- (Exam Topic 5)
You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table.
You need to identify the extent of the data skew in Table1. What should you do in Synapse Studio?

A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
B. Connect to the built-in pool and run DBCC CHECKALLOC.
C. Connect to Pool1 and run DBCC CHECKALLOC.
D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Answer: D

Explanation:
Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data. Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet

NEW QUESTION 131
- (Exam Topic 5)
You have the following Azure Resource Manager template.

```
...
  "variable": {
    "serverName": "azsqldbserver0001"
  },
  "resources": [
    {
      "name": "[variables('serverName')]",
      "type": "Microsoft.Sql/servers",
      "apiVersion": "2019-06-01-preview",
      "location": "[parameters('location')]",
      "properties": {
        "administratorLogin": "[parameters('administratorLogin')]",
        "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
        "version": "12.0"
      },
      "resources": [
        {
          "name": "[concat(variables('serverName'),'/',parameters('databaseName'))]",
          "type": "Microsoft.Sql/servers/databases",
          "apiVersion": "2020-08-01-preview",
          "location": "[parameters('location')]",
          "kind": "v12.0",
          "sku": {
            "name": "Standard",
            "tier": "Standard",
            "capacity": 10
          },
          "dependsOn": [
            "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
          ],
          "properties": {
          },
          "resources": [
          ]
        }
      ]
    }
  ],
...
```

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

Statements	Yes	No
The template deploys a serverless Azure SQL database.	O	O
The template deploys a database to an Azure SQL Database managed instance.	O	O
The pricing tier of the database deployment is based on DTUs.	O	O

A. Mastered
B. Not Mastered

Answer: A

Explanation:
A screenshot of a computer Description automatically generated with low confidence
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/purchasing-models https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart

NEW QUESTION 134
- (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 136
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.
You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.
Solution: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.
Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

NEW QUESTION 137
- (Exam Topic 5)
You are building an Azure Stream Analytics job to retrieve game data.
You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.
How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

```
SELECT ▼ as HighestScore
       Collect(Score)
       CollectTop(1)OVER(ORDER BY Score Desc)
       Game, MAX(Score)
       TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

FROM input TIMESTAMP BY CreatedAt

GROUP BY ▼
         Game
         Hopping(minute, 5)
         Tumbling(minute, 5)
         Windows(TumblingWindow(minute, 5), Hopping(minute, 5))
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)
TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.
Analytic Function Syntax:
TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT
DURATION clause> [<WHEN clause>])
Box 2: Tumbling(minute 5)
Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The

key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

Tell me the count of Tweets per time zone every 10 seconds


```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```
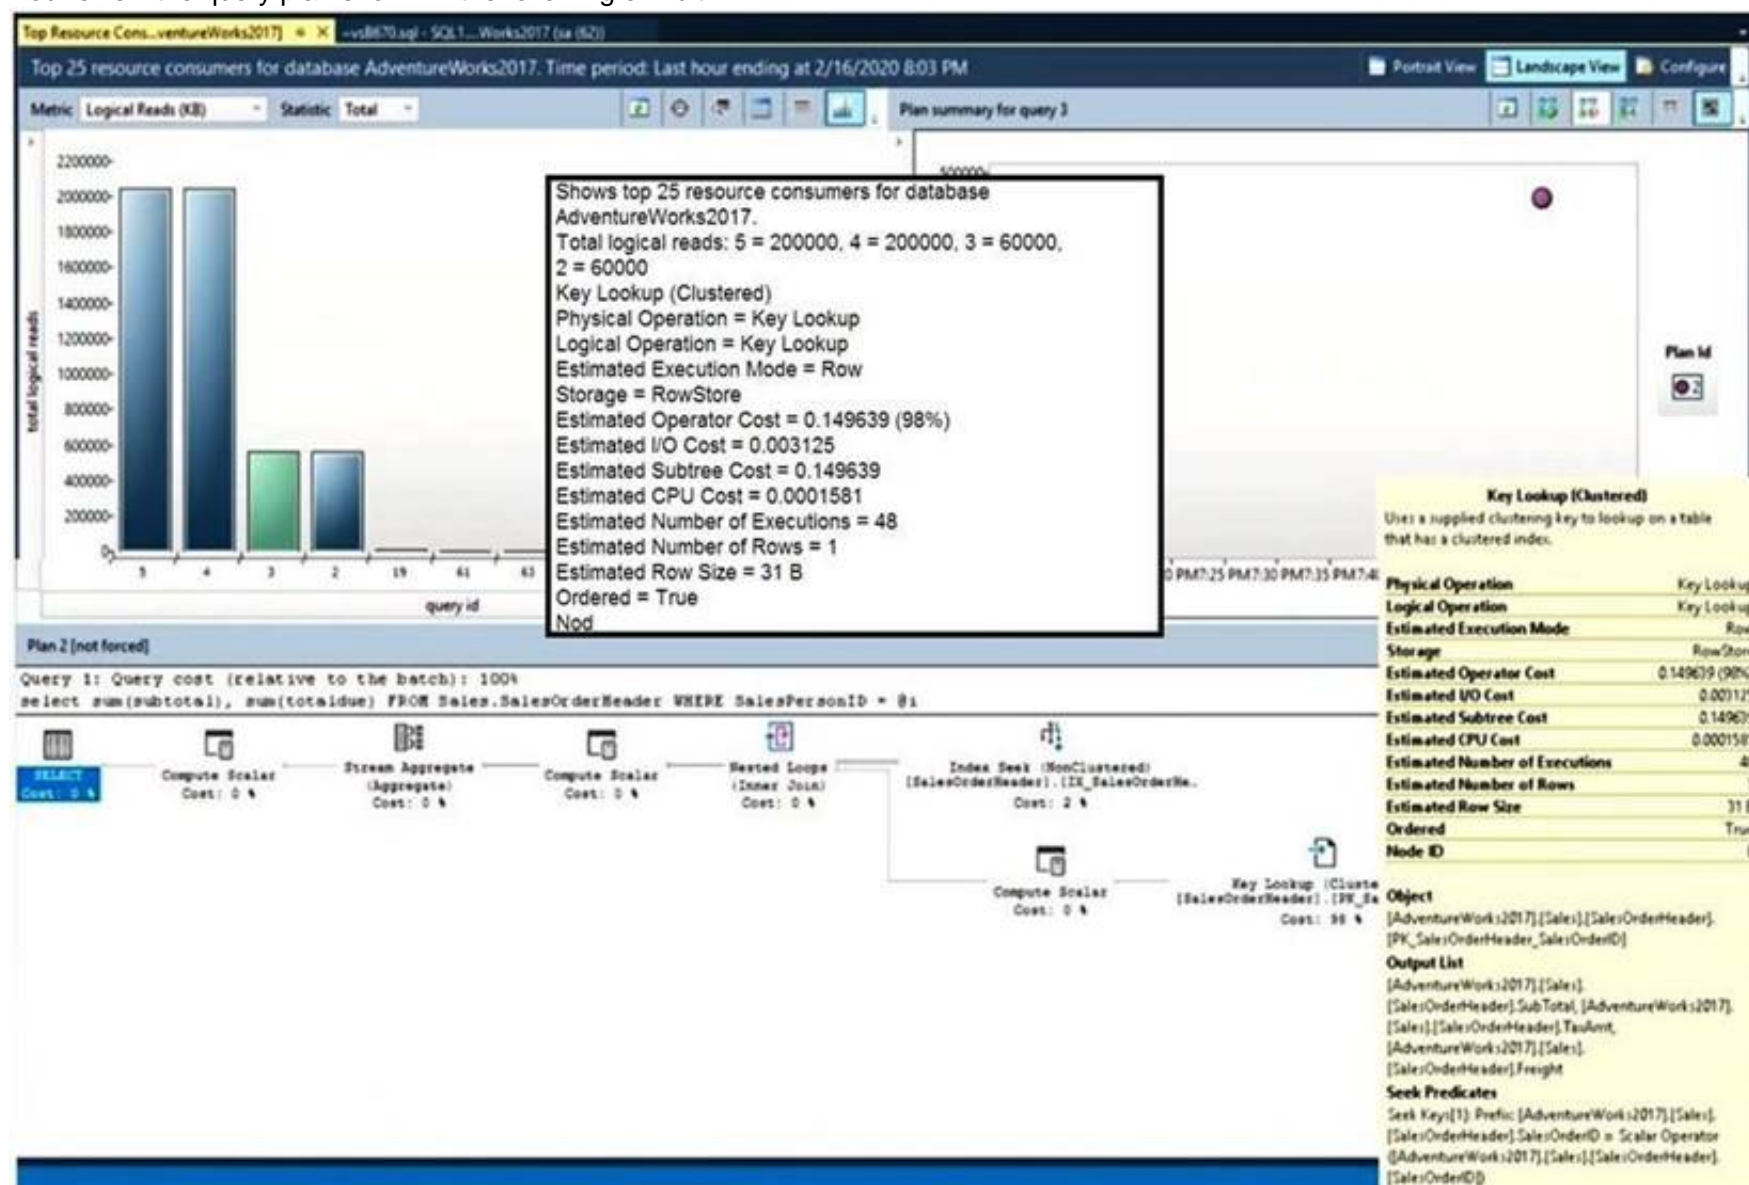
Reference:
https://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-fun

NEW QUESTION 139
- (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 143
- (Exam Topic 5)
You have a database on a SQL Server on Azure Virtual Machines instance.
The current state of Query Store for the database is shown in the following exhibit.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text Description automatically generated

NEW QUESTION 145
- (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 148
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to display the estimated execution plan of a query by using the query editor in the Azure portal.
What should you do first?

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

Answer: A

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

NEW QUESTION 153
- (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.

Actions	Answer Area
Create a workload group.	
Create a user-defined classifier function.	
Modify Resource Governor.	
Create a contained database user.	
Create a resource pool.	

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 156
- (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 159
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL resource that will support cross database queries by using an Azure Resource Manager (ARM) template.
How should you complete the ARM template? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Answer Area

```
"resources": [
    ...
    "type":        [ Microsoft.Sql/servers          ▼ ]
                     Microsoft.Sql/servers/databases
                     Microsoft.Sql/managedInstances

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 164
- (Exam Topic 5)
You have an Azure SQL database named DB1.
You need to encrypt DB1. The solution must meet the following requirements;
• Encrypt data in motion.
• Support comparison operators.
• Provide randomized encryption.
What should you include in the solution?

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

Answer: A

NEW QUESTION 166
- (Exam Topic 5)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:

➤ During normal operations, provide at least two readable copies of Sales.

➤ Ensure that Sales remains available if a datacenter fails.
Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage

(locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

NEW QUESTION 168
- (Exam Topic 5)
You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

≫ Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.

≫ Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: B

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

NEW QUESTION 173
- (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.

Configuration setting:
```
DATA_FLUSH_INTERVAL_SECONDS
INTERVAL_LENGTH_MINUTES
MAX_PLANS_PER_QUERY
QUERY_CAPTURE_MODE
```

Value:
```
1
60
CUSTOM
ON
```

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 174
- (Exam Topic 5)
You have an Azure Databricks workspace named workspace1 in the Standard pricing tier. Workspace1 contains an all-purpose cluster named cluster1.
You need to reduce the time it takes for cluster1 to start and scale up. The solution must minimize costs. What should you do first?

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

Answer: C

Explanation:
You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.

Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.
Reference:
https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html

NEW QUESTION 179
- (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 183
- (Exam Topic 5)
You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the
Exhibit tab.)

Configure policies
SQL server

Point in Time Restore Configuration

| Configure PiTR backup retention ⌄ | Days |

Long-term Retention Configurations

☑ Weekly LTR Backups ⓘ

How long would you like weekly backups to be kept?

| 6 ✓ | Week(s) ⌄ |

☑ Monthly LTR Backups ⓘ

How long would you like the first backup of each month to be kept?

| 12 ✓ | Month(s) ⌄ |

☑ Yearly LTR Backups ⓘ

Which weekly backup of the year would you like to retain?

| Week 2 ⌄ |

How long would you like this annual backup to be kept?

| 10 ✓ | Year(s) ⌄ |

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.

The backup saved to long-term retention on January 4, 2020, will be retained for [▼]
| 6 weeks |
| 12 months |
| 10 years |

The backup saved to long-term retention on January 11, 2020 will be retained for [▼]
| 6 weeks |
| 12 months |
| 10 years |

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 185
- (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 188
- (Exam Topic 5)
You have an Azure SQL managed instance that hosts multiple databases.
You need to configure alerts for each database based on the diagnostics telemetry of the database. What should you use?

A. Azure SQL Analytics alerts based on metrics
B. SQL Health Check alerts based on diagnostics logs
C. SQL Health Check alerts based on metrics
D. Azure SQL Analytics alerts based on diagnostics logs

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-expo

NEW QUESTION 191
- (Exam Topic 5)
You have a burstable Azure virtual machine named VMI that hosts an instance of Microsoft SQL Server. You need to attach an Azure ultra disk to VMI. The solution must minimize downtime on VMI.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions		Answer Area
Attach the ultra disk.		
Stop and deallocate VM1.		
Set Enable Ultra disk compatibility to **Yes**.	❯	
Resize VM1.	❮	
Start VM1.		

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

Attach the ultra disk.
Stop and deallocate VM1.
Set Enable Ultra disk compatibility to **Yes**.
Resize VM1.
Start VM1.

Answer Area

Stop and deallocate VM1.
Attach the ultra disk.
Set Enable Ultra disk compatibility to **Yes**.
Resize VM1.
Start VM1.

NEW QUESTION 192
- (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 197
- (Exam Topic 5)
From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.
The data contains the following columns:

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

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

EventCategory:

| DimChannel |
| DimDate |
| DimEvent |
| FactEvents |

ChannelGrouping:

| DimChannel |
| DimDate |
| DimEvent |
| FactEvents |

TotalEvents:

| DimChannel |
| DimDate |
| DimEvent |
| FactEvents |

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 201
- (Exam Topic 5)
You have an Azure subscription.
You need to deploy an Azure SQL managed instance that meets the following requirements:
•Optimize latency.
•Maximize the memory-to-vCore ratio.
Which service tier and hardware generation should you use? To answer, select the apocopate options in the answer area.
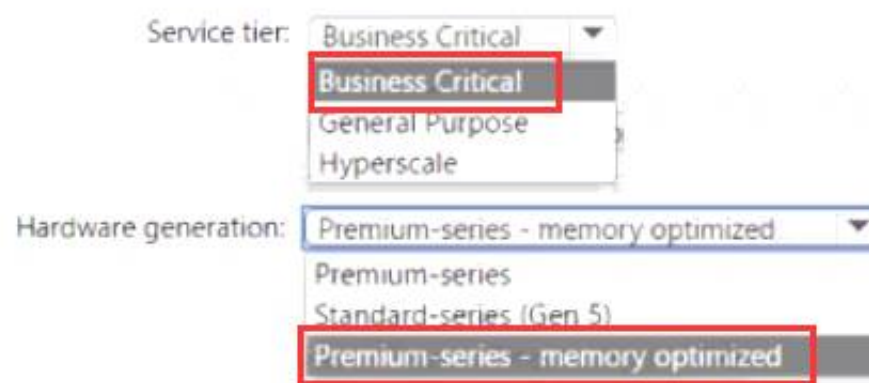NOTE: Each correct selection is worth one point.

Answer Area

Service tier: Business Critical

| Business Critical |
| General Purpose |
| Hyperscale |

Hardware generation: Premium-series - memory optimized

| Premium-series |
| Standard-series (Gen 5) |
| Premium-series - memory optimized |

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Service tier: Business Critical ▼

Business Critical
General Purpose
Hyperscale

Hardware generation: Premium-series - memory optimized ▼

Premium-series
Standard-series (Gen 5)
Premium-series - memory optimized

NEW QUESTION 206
- (Exam Topic 5)
You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.
Data to be loaded is identified by a column named LastUpdatedDate in the source table. You plan to execute the pipeline every four hours.
You need to ensure that the pipeline execution meets the following requirements:
Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.
Which type of trigger should you use?

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

Answer: A

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

NEW QUESTION 211
- (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 212
- (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 213
- (Exam Topic 3)
Which windowing function should you use to perform the streaming aggregation of the sales data?
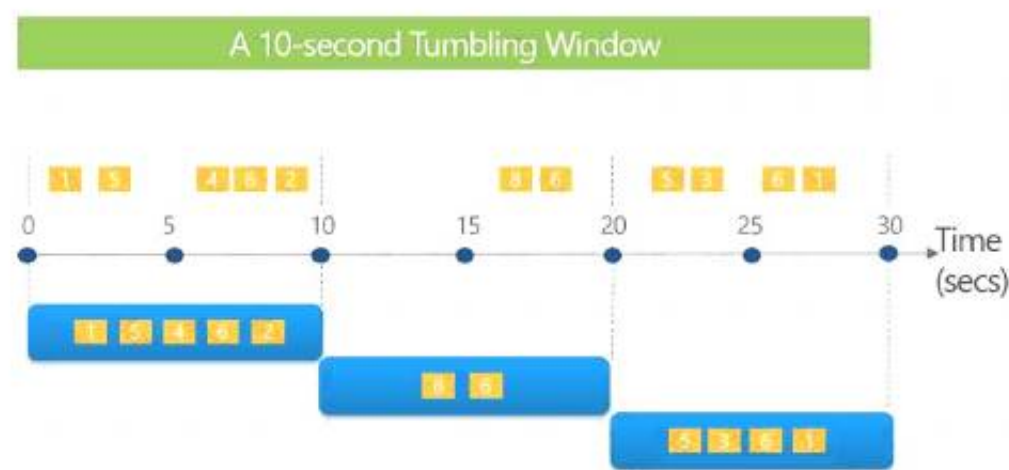
A. Sliding
B. Hopping

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


```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

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

NEW QUESTION 218
- (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 220
- (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 223
- (Exam Topic 2)
What should you use to migrate the PostgreSQL database?

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

Answer: C

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

NEW QUESTION 224
- (Exam Topic 1)
You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

Purchasing model:
Azure virtual machine reserved instances
DTU
vCore

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: DTU
Scenario:
≫ The 30 new databases must scale automatically.
≫ Once all requirements are met, minimize costs whenever possible.
You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.
In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.
Box 2: An Azure SQL database elastic pool
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

NEW QUESTION 226
- (Exam Topic 1)
What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

A. Configure SalesSQLDb1 as writable.
B. Update the connection strings of SalesSQLDb1App1.
C. Update the firewall rules of SalesSQLDb1.

D. Update the users in SalesSQLDb1.

Answer: B

Explanation:
Scenario: SalesSQLDb1 uses database firewall rules and contained database users.

NEW QUESTION 229
- (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.

Actions	Answer Area
Create a user in the master database	
Modify the Azure SQL server administrator account	
Create contained database users	
Create an Azure AD administrator for the logical server	
Connect to the databases by using an Azure AD account	
Enable the contained database authentication option	

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 231
- (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.

Collect metrics from: ▼

| The database only |
| The elastic pool and the database |
| The elastic pool only |
| The server, the elastic pool, and the database |

Stream metrics to: ▼

| Azure Event Hubs |
| Azure Log Analytics |
| Azure Storage |

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 234
- (Exam Topic 1)
You need to identify the cause of the performance issues on SalesSQLDb1.
Which two dynamic management views should you use? Each correct answer presents part of the solution.
NOTE:Each correct selection is worth one point.

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

Answer: AE

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

NEW QUESTION 237
- (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.

Quorum model:
- Cloud witness
- Disk witness
- File share witness

Azure resource for the availability group listener:
- Azure Application Gateway
- Azure Basic Load Balancer

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 241
......

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