

Microsoft

Exam Questions DP-300

Administering Relational Databases on Microsoft Azure (beta)



NEW QUESTION 1

- (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 2

- (Exam Topic 5)

You have an Azure SQL database.

You discover that the plan cache is full of compiled plans that were used only once.

You run theselect * from sys.database_scoped_configurationsTransact-SQL command and receive the results shown in the following table.

configuration_id	name	value	is_value_default
1	LEGACY_CARDINALITY_ESTIMATION	0	1
2	QUERY_OPTIMIZER_HOTFIXES	0	1
3	OPTIMIZE_FOR_AD_HOC_WORKLOADS	0	1
4	ACCELERATED_PLAN_FORCING	1	1

You need relieve the memory pressure. What should you configure?

- A. LEGACY_CARDINALITY_ESTIMATION
- B. QUERY_OPTIMIZER_HOTFIXES
- C. OPTIMIZE_FOR_AD_HOC_WORKLOADS
- D. ACCELERATED_PLAN_FORCING

Answer: C

Explanation:

OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }

Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql>

NEW QUESTION 3

- (Exam Topic 5)

You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux.

You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

To prepare for the migration:

To perform the migration, use:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To prepare for the migration: ▼

To perform the migration, use: ▼

NEW QUESTION 4

- (Exam Topic 5)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query `sys.dm_exec_requests` and discover that the wait type is `PAGELATCH_UP` and the wait resource is `2:3:905856`.

You need to improve system performance.

Solution: You reduce the use of table variables and temporary tables. 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 5

- (Exam Topic 5)

You have an Azure Active Directory (Azure AD) tenant named `contoso.com` that contains a user named `user1@contoso.com` and an Azure SQL managed instance named `SQLMI1`.

You need to ensure that `user1@contoso.com` can create logins in `SQLMI1` that map to Azure AD service principals.

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

← ↑

→ ↓

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with medium confidence

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/aad-security-configure-tutorial>

NEW QUESTION 6

- (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. Run the `DBCC SQLPERF` command.
- B. Run the `dbcc dbreindex` command.
- C. Modify the automatic tuning settings for `db1`.

D. Modify the Query Store settings for db1.

Answer: C

Explanation:

Reference:

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

NEW QUESTION 7

- (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 is NOT 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 is 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

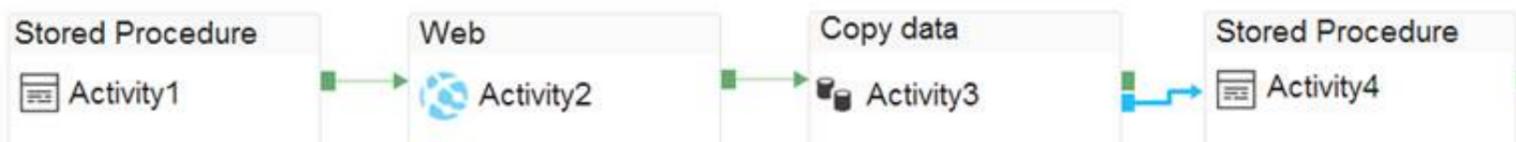
Original (raw) data that must be preserved, even after it has been processed into final usable form Compliance and archival data that needs to be stored for a long time and is hardly ever accessed Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

NEW QUESTION 8

- (Exam Topic 5)

You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

- > Operator: Greater than
- > Aggregation type: Total
- > Threshold value: 2
- > Aggregation granularity (Period): 5 minutes
- > Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

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
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: No

Just one failure within the 5-minute interval. Box 2: No

Just two failures within the 5-minute interval. Box 3: No

Just two failures within the 5-minute interval. Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

NEW QUESTION 9

- (Exam Topic 5)

You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1. You plan to implement Azure SQL Data Sync for DB1. Which isolation level should you configure?

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

Answer: B

Explanation:

Reference:

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

NEW QUESTION 10

- (Exam Topic 5)

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.

You need to ensure that an automatic email notification is sent once the job completes. What should you include in the solution?

- A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
- B. From SQL Server Management Studio (SSMS), runsp_set_sqlagent_properties
- C. From SQL Server Management Studio (SSMS), create a Database Mail profile
- D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

Answer: C

Explanation:

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail. Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:

- > In Object Explorer, expand a SQL Server instance.
- > Right-click SQL Server Agent, and then click Properties.
- > Click Alert System.
- > Select Enable Mail Profile.
- > In the Mail system list, select Database Mail.
- > In the Mail profile list, select a mail profile for Database Mail.
- >

Restart SQL Server Agent.

Note: Prerequisites include:

- Enable Database Mail.
- Create a Database Mail account for the SQL Server Agent service account to use.
- Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.
- Set the profile as the default profile for the msdb database. Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d>

NEW QUESTION 10

- (Exam Topic 5)

You have an Azure subscription that contains an instance of SQL Server on an Azure virtual machine named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.

You need to ensure that User1 can create a scheduled task to perform a full backup of DB1. The solution must use the principle of least privilege.

Which built-in database role should you assign to User1?

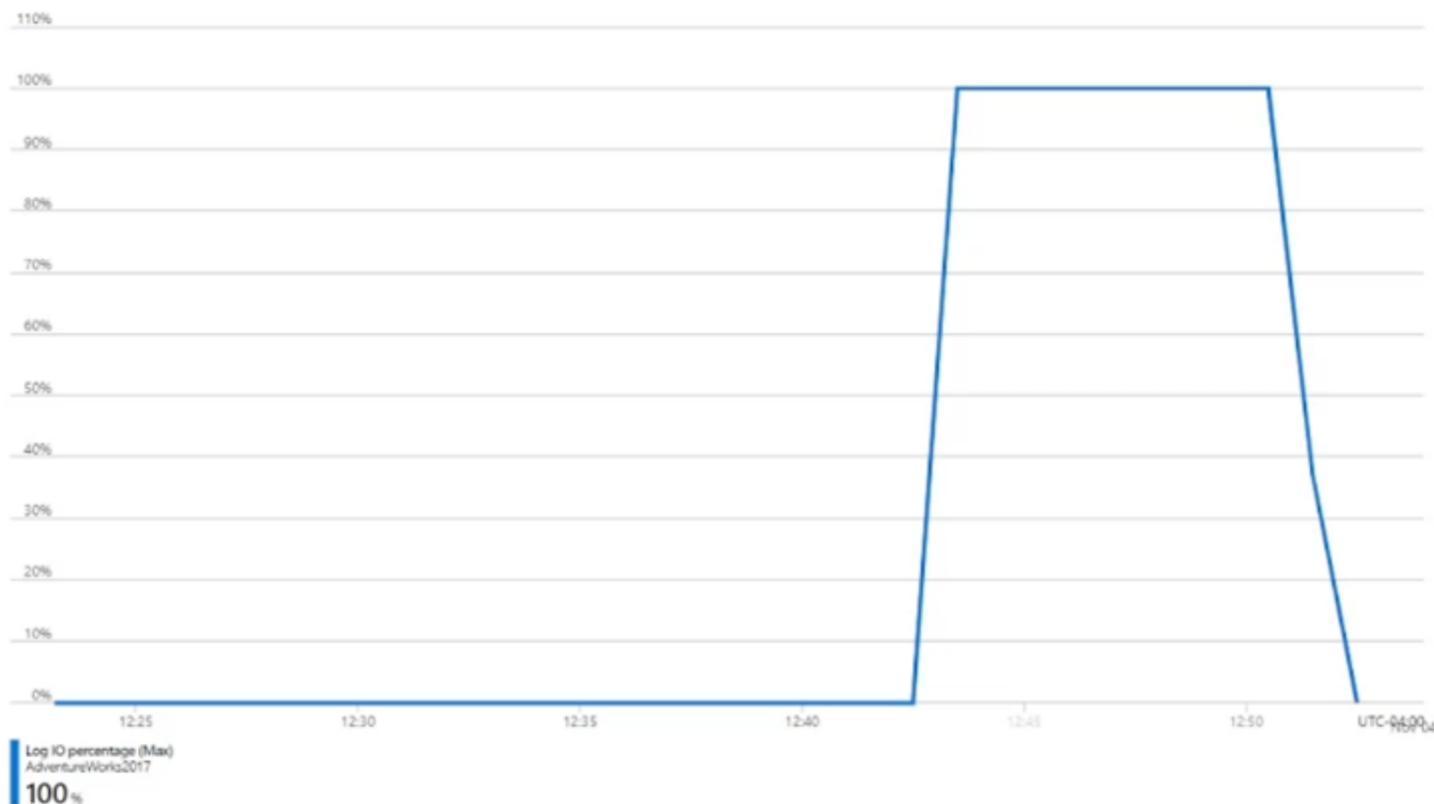
- A. SQLAgentReaderRole
- B. db.owner
- C. SQLAgentOperatorRole
- D. SQLAgentUserRole

Answer: C

NEW QUESTION 12

- (Exam Topic 5)

You have an Azure SQL database named DB1 in the General Purpose service tier. The performance metrics for DB1 are shown in the following exhibit.



You need to reduce the Log 10 percentage. The solution must minimize costs. What should you do?

- A. Increase the number of vCores.
- B. Change Recoverymodel to Simple.
- C. Perform a checkpoint operation.
- D. Change Service tier to Business Critical.

Answer: D

NEW QUESTION 14

- (Exam Topic 5)

You plan to perform batch processing in Azure Databricks once daily. Which type of Databricks cluster should you use?

- A. automated
- B. interactive
- C. High Concurrency

Answer: A

Explanation:

Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.

The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/clusters>

NEW QUESTION 17

- (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 >

▼
DATEADD
DATEDIFF
DATEPART
TODATETIMEOFFSET

(day, -7, GETDATE())

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 18

- (Exam Topic 5)

You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in five-minute intervals and report only events that arrive during the interval.

The output will be sent to a Delta Lake table. Which output mode should you use?

- A. complete
- B. append
- C. update

Answer: A

Explanation:

Complete mode: You can use Structured Streaming to replace the entire table with every batch.

Reference:

<https://docs.databricks.com/delta/delta-streaming.html>

NEW QUESTION 20

- (Exam Topic 5)

You have an Azure SQL database named DB1.

You have a table name Table1 that has 20 columns of type CHAR(400). Row compression for Table1 is enabled.

During a database audit, you discover that none of the fields contain more than 150 characters. You need to ensure that you can apply page compression to Table1.

What should you do?

- A. Configure the columns as sparse.
- B. Change the column type to nvarchar(MAX).
- C. Change the column type to varchar(MAX).
- D. Change the column type to varchar(200).

Answer: D

Explanation:

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/> <https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/>

NEW QUESTION 24

- (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 FILE LOG		DISK = N'\\NAS01\SSQLBackups\DB1.bak'; TAPE = N'\\Tape0' 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 27

- (Exam Topic 5)

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.

You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

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

Answer: C

Explanation:

Use the Query Store Page in SQL Server Management Studio.

Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query.

This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

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

NEW QUESTION 31

- (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 33

- (Exam Topic 5)

You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.

Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.

You set the Read/Write failover policy to Automatic.

What are two results of the configuration? Each correct answer presents a complete solution.

NOTE:Each correct selection is worth one point.

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

E. In the event of an outage, the minimum delay for the databases to fail over in the primary instance will be one hour.

Answer: AE

Explanation:

A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.
E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

NEW QUESTION 35

- (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 38

- (Exam Topic 5)

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- > Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- > Restrict network connectivity to SqlSrv1. What should you create on VNet1?

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

Answer: C

Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- > Azure Storage
- > Azure Cosmos DB
- > Azure SQL Database
- > Your own service using a Private Link Service. Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview>

NEW QUESTION 41

- (Exam Topic 5)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You restore Database1 from Server1 to the Server2 by using theRESTORETransact-SQL command and theREPLACEoption.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

With the REPLACE option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

NEW QUESTION 45

- (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 47

- (Exam Topic 5)

You have the following Azure Data Factory pipelines:

- > Ingest Data from System1
- > Ingest Data from System2
- > Populate Dimensions
- > Populate Facts

Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2. Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours.

What should you do to schedule the pipelines for execution?

- A. Add a schedule trigger to all four pipelines.
- B. Add an event trigger to all four pipelines.
- C. Create a parent pipeline that contains the four pipelines and use an event trigger.
- D. Create a parent pipeline that contains the four pipelines and use a schedule trigger.

Answer: D

Explanation:

Reference:

<https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/>

NEW QUESTION 51

- (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

Answer Area

- 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];



- 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 55

- (Exam Topic 5)

You have an Azure Synapse Analytics dedicated SQL pool.

You run PDW_SHOWSPACEUSED('dbo.FactInternetSales'); and get the results shown in the following table.

ROWS	RESERVED_SPACE	DATA_SPACE	INDEX_SPACE	UNUSED_SPACE	PDW_NODE_ID	DISTRIBUTION_ID
694	2776	616	48	2112	1	1
407	2704	576	48	2080	1	2
53	2376	512	16	1848	1	3
58	2376	512	16	1848	1	4
168	2632	528	32	2072	1	5
195	2696	536	32	2128	1	6
5995	3464	1424	32	2008	1	7
0	2232	496	0	1736	1	8
264	2576	544	40	1992	1	9
3008	3016	960	32	2024	1	10
...
1550	2832	752	48	2032	1	50
1238	2832	696	40	2096	1	51
192	2632	528	32	2072	1	52
1127	2768	680	48	2040	1	53
1244	3032	704	64	2264	1	54
409	2632	568	32	2032	1	55
0	2232	496	0	1736	1	56
1437	2832	728	40	2064	1	57
0	2232	496	0	1736	1	58
384	2632	560	32	2040	1	59
225	2768	544	40	2184	1	60

Which statement accurately describes the dbo.FactInternetSales table?

- A. The table contains less than 10,000 rows.
- B. All distributions contain data.
- C. The table uses round-robin distribution
- D. The table is skewed.

Answer: D

Explanation:

The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.

Note: SHOWSPACEUSED displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database.

This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.

ROUND_ROBIN distributed tables should not be skewed. Data is distributed evenly across the nodes by design.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribu> <https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distrib>

NEW QUESTION 59

- (Exam Topic 5)

You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1. You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. From the properties of DB1. configure geo-replication.
- B. From the properties of Server1 add a failover group.
- C. Create a new Azure SQL Database server named Server2.
- D. From the properties of Server1 configure retention for DB1
- E. Create a new Azure SQL Database instance named DB2.

Answer: BC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell> <https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azur>

NEW QUESTION 61

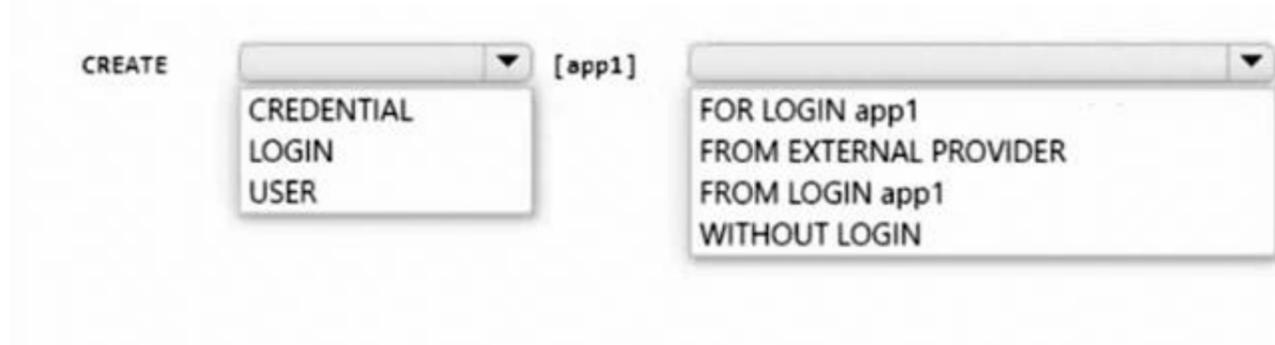
- (Exam Topic 5)

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named appl.

You need to ensure that appl can access db1. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 63

- (Exam Topic 5)

You are designing a star schema for a dataset that contains records of online orders. Each record includes an order date, an order due date, and an order ship date.

You need to ensure that the design provides the fastest query times of the records when querying for arbitrary date ranges and aggregating by fiscal calendar attributes.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Create a date dimension table that has a DateTime key.
- B. Create a date dimension table that has an integer key in the format of YYYYMMDD.
- C. Use built-in SQL functions to extract date attributes.
- D. Use integer columns for the date fields.
- E. Use DateTime columns for the date fields.

Answer: BD

Explanation:

Reference:

https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-in-a

NEW QUESTION 66

- (Exam Topic 5)

You have 20 Azure SQL databases provisioned by using the vCore purchasing model. You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases

- B. geo-replication support
- C. number of concurrently peaking databases * peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases * average CPU utilization per database

Answer: ACE

Explanation:

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model: MAX(<Total number of DBs X average vCore utilization per DB>, <Number of concurrently peaking DBs X Peak vCore utilization per DB>)

A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

NEW QUESTION 69

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1. You need to move a SQL Server Agent job from DB1 to SQLMI1. Which job attribute is unsupported in SQLMI1?

- A. log to table
- B. email notifications
- C. schedules
- D. output files

Answer: D

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"}
```

<ul style="list-style-type: none"> - assignidentity - federatedclientID - keyid 	<ul style="list-style-type: none"> - assignidentity - federatedclientID - keyid
--	--

```
"https://db1.vault.azure.net/keys/dbkey/01234234512345678901234561823942"
```

- 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 SQL Server on an Azure virtual machine.

You need to use Policy-Based Management in Microsoft SQL Server to identify stored procedures that do not comply with your naming conventions.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Export a built-in policy.
- Create a custom policy based on a condition.
- Create a custom condition based on a built-in facet.
- View the policy history.
- Import a policy file.
- Run a policy evaluation.

Answer Area

⏪	⏩
⏴	⏵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Reference:

<https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-bas>

NEW QUESTION 77

- (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 81

- (Exam Topic 5)

You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2 . SQLMI2 contains a database named DB1 and a user named User1.

User1 drops DB1.

You need to perform a point-in-time restore of DB1 to SQLMI2.

- A. Azure CLI
- B. Transact-SQL
- C. The Azure portal
- D. Azure PowerShell

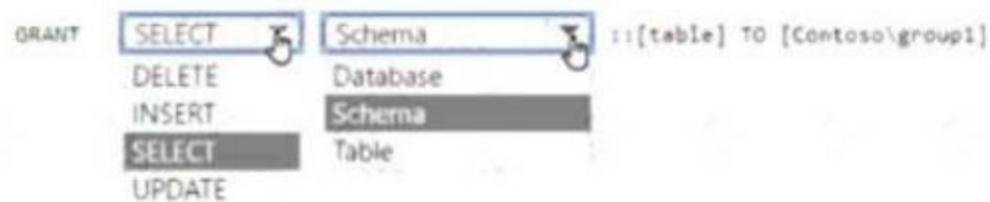
Answer: C

NEW QUESTION 83

- (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 in 061. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 87

- (Exam Topic 5)

You plan to deploy an app that includes an Azure SQL database and an Azure web app. The app has the following requirements:

- > The web app must be hosted on an Azure virtual network.
- > The Azure SQL database must be assigned a private IP address.
- > The Azure SQL database must allow connections only from the virtual network. You need to recommend a solution that meets the requirements.

What should you include in the recommendation?

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

Answer: A

Explanation:

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

NEW QUESTION 89

- (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 93

- (Exam Topic 5)

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

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

Answer: A

Explanation:

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.

Reference:
<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview>

NEW QUESTION 97

- (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 98

- (Exam Topic 5)

You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1. You need to configure SQL1 to use mixed mode authentication. Which procedure should you run?

- A. sp_addremotelogin
- B. xp_instance_regwrite
- C. sp_cncharge_users_login
- D. xp_grant_login

Answer: B

NEW QUESTION 99

- (Exam Topic 5)

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

- A. Restore with Recovery
- B. Restore with Norecovery
- C. Restore with Standby

Answer: B

Explanation:

Prepare a secondary database for an Always On availability group requires two steps:

* 1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY

* 2. Join the restored database to the availability group. Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server>

NEW QUESTION 100

- (Exam Topic 5)

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

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

Answer: D

Explanation:

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

NEW QUESTION 104

- (Exam Topic 5) You have an Azure SQL database. You identify a long running query.

You need to identify which operation in the query is causing the performance issue.

What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

- A. Live Query Statistics
- B. an estimated execution plan
- C. an actual execution plan
- D. Client Statistics

Answer: A

Explanation:

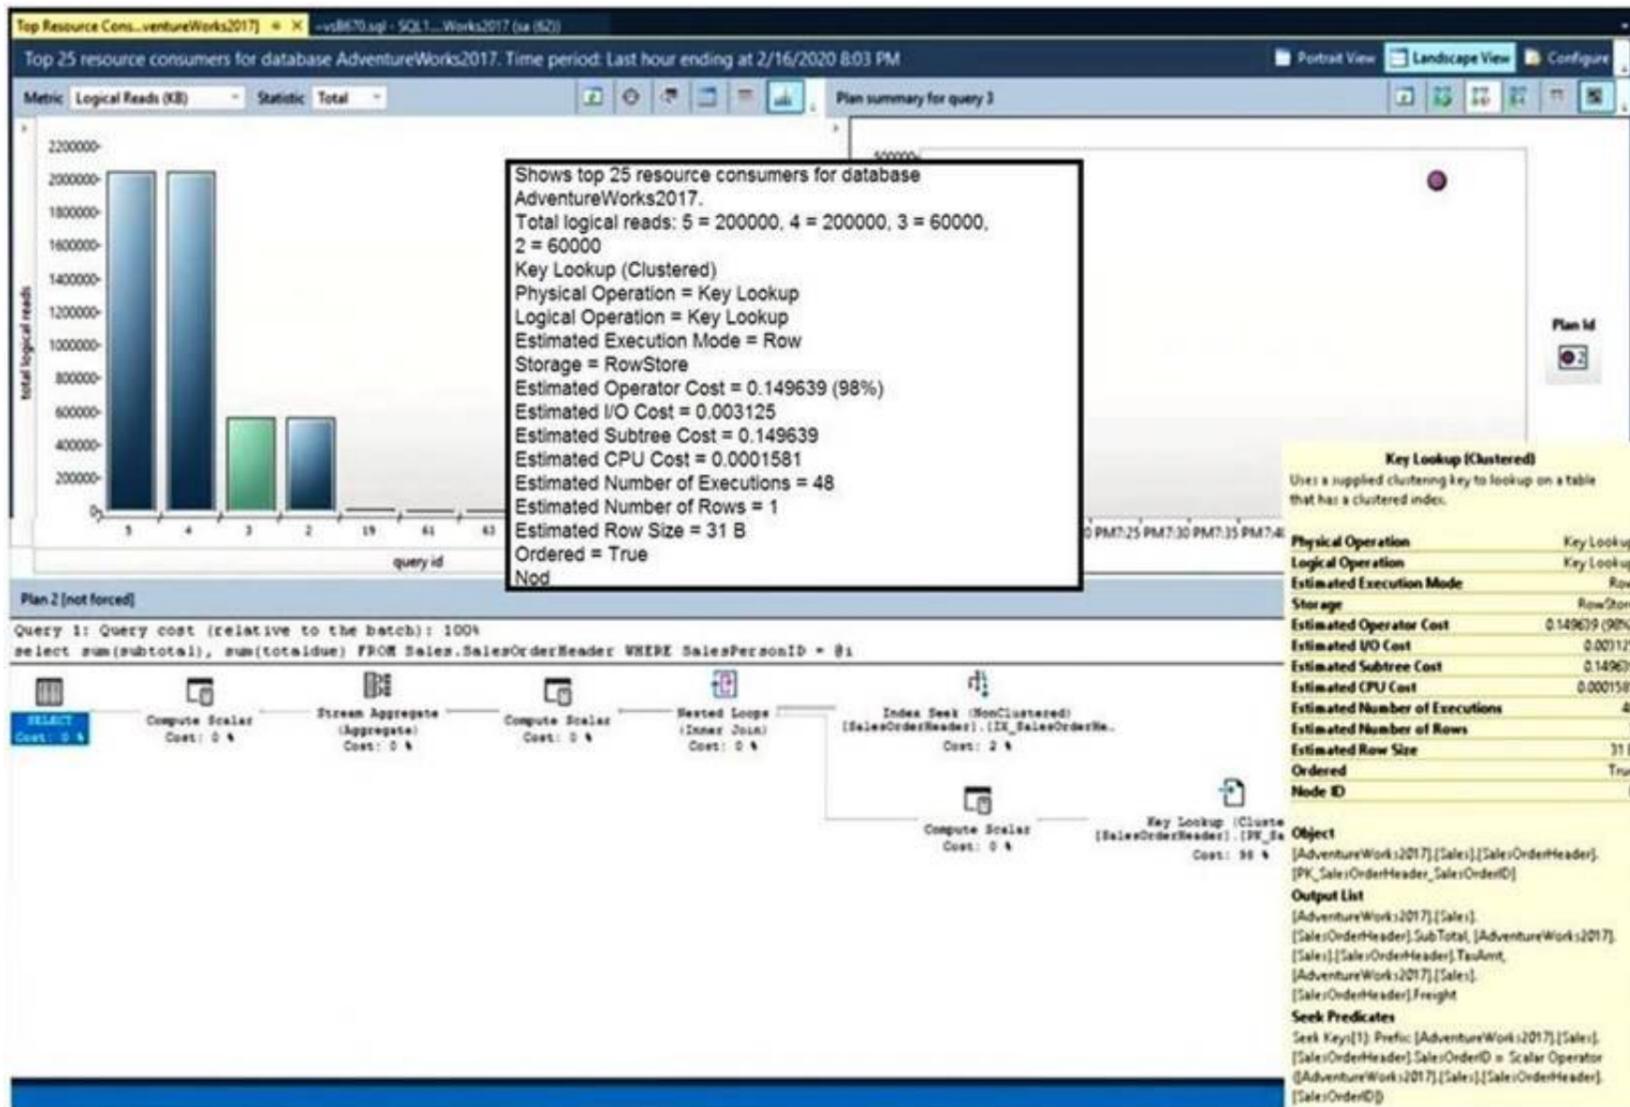
<https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/>

NEW QUESTION 107

- (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

- 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.

Yes	No
<input type="radio"/>	<input type="radio"/>
<input type="radio"/>	<input type="radio"/>
<input type="radio"/>	<input type="radio"/>

- 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 112

- (Exam Topic 5)

You have a SQL pool in Azure Synapse that contains a table named dbo.Customers. The table contains a column name Email. You need to prevent nonadministrative users from seeing the full email addresses in the Email column. The users must see values in a format of aXXX@XXXX.com instead. What should you do?

- A. From the Azure portal, set a mask on the Email column.
- B. From the Azure portal, set a sensitivity classification of Confidential for the Email column.
- C. From Microsoft SQL Server Management Studio, set an email mask on the Email column.
- D. From Microsoft SQL Server Management Studio, grant the SELECT permission to the users for all the columns in the dbo.Customers table except Email.

Answer: B

Explanation:

The Email masking method, which exposes the first letter and replaces the domain with XXX.com using a constant string prefix in the form of an email address. Example: aXX@XXXX.com

NEW QUESTION 114

- (Exam Topic 5)

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

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

Answer: A

Explanation:

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

NEW QUESTION 119

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

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

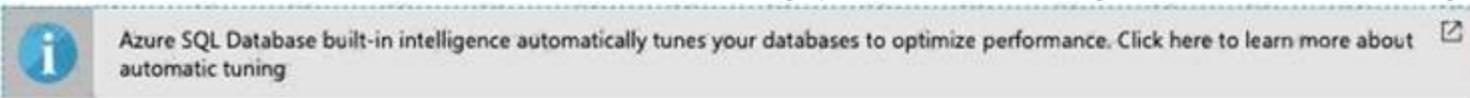
Answer: A

Explanation:

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

NEW QUESTION 124

- (Exam Topic 5)
 You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.



Inherit from: ⓘ

Server
 Azure defaults
 Don't inherit

ⓘ The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options ⓘ

OPTION	DESIRED STATE	CURRENT STATE
FORCE PLAN	<input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT	ON Auto-configured by Azure
CREATE INDEX	<input type="radio"/> ON <input type="radio"/> OFF <input checked="" type="radio"/> INHERIT	ON Auto-configured by Azure
DROP INDEX	<input checked="" type="radio"/> ON <input type="radio"/> OFF <input type="radio"/> INHERIT	ON Forced by user

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
Nonclustered indexes will be added to tables to improve performance.	<input type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

We see: Tuning option: Create index ON

CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.

Box 2: No

Box 3: Yes

FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

NEW QUESTION 127

- (Exam Topic 5)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

- > During normal operations, provide at least two readable copies of Sales.
- > Ensure that Sales remains available if a datacenter fails.

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

- A. Yes
- B. No

Answer: B

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

NEW QUESTION 129

- (Exam Topic 5)

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Answer: C

Explanation:

A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:

What resources the client may access.

What permissions they have to those resources. How long the SAS is valid.

Note: Data Lake Storage Gen2 supports the following authorization mechanisms:

- > Shared Key authorization
- > Shared access signature (SAS) authorization
- >

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

Reference:

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

NEW QUESTION 132

- (Exam Topic 5)

You have an Azure SQL database named sqldb1.

You need to minimize the amount of space by the data and log files of sqldb1. What should you run?

- A. DBCC SHRINKDATABASE
- B. sp_clean_db_free_space
- C. sp_clean_db_file_free_space
- D. DBCC SHRINKFILE

Answer: D

Explanation:

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database. Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

NEW QUESTION 135

- (Exam Topic 5)

You have an Azure SQL logical server. You run the following script.

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

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

NOTE: Each correct selection is worth one point

Statements	Yes	No
The <code>orders</code> table will allow only rows to be inserted.	<input type="radio"/>	<input type="radio"/>
To create additional tables in the Sales database, the <code>LEDGER = ON</code> parameter must be used.	<input type="radio"/>	<input type="radio"/>
To ensure that a timestamp is added to each row in the <code>orders</code> table, the <code>GENERATED ALWAYS</code>	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Yes

No No

NEW QUESTION 137

- (Exam Topic 5)

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- > Clustered columnstore indexes
- > Automatic tuning
- > Change tracking
- > PolyBase

You plan to migrate DB1 to an Azure SQL database.
 What feature should be removed or replaced before DB1 can be migrated?

- A. Clustered columnstore indexes
- B. PolyBase
- C. Change tracking
- D. Automatic tuning

Answer: B

Explanation:

This table lists the key features for PolyBase and the products in which they're available.

Feature	SQL Server (Beginning with 2016)	Azure SQL Database	Azure Synapse Analytics	Parallel Data Warehouse
Query Hadoop data with Transact-SQL	Yes	No	No	Yes
Import data from Hadoop	Yes	No	No	Yes
Export data to Hadoop	Yes	No	No	Yes
Query, import from, export to Azure HDInsight	No	No	No	No
Push down query computations to Hadoop	Yes	No	No	Yes
Import data from Azure Blob storage	Yes	Yes	Yes	Yes
Export data to Azure Blob storage	Yes	No	Yes	Yes
Import data from Azure Data Lake Store	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft BI tools	Yes	No	Yes	Yes

Reference:
<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary>

NEW QUESTION 142

- (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 146

- (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 150

- (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 151

- (Exam Topic 5)

You have an Azure SQL managed instance.

You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 155

- (Exam Topic 5)

You have a on-premises Microsoft SQL Server named SQL1 that hosts five databases.

You need to migrate the databases to an Azure SQL managed instance. The solution must minimize downtime and prevent data loss. What should you use?

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

Answer: C

NEW QUESTION 157

- (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 159

- (Exam Topic 5)

You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:

- Support real-time data replication to a different geographic region.
- Use Azure as a disaster recovery target.
- Minimize costs and administrative effort.

What should you include in the recommendation?

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

Answer: D

NEW QUESTION 161

- (Exam Topic 5)

You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1. You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.

What should you include in the recommendation?

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

Answer: B

NEW QUESTION 163

- (Exam Topic 5)

You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.

You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.

Which resource type should you use to create the alert?

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

Answer: C

Explanation:

There are resource types related to application code, compute infrastructure, networking, storage + databases. You can deploy up to 800 instances of a resource type in each resource group.

Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types>

NEW QUESTION 164

- (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 169

- (Exam Topic 5)

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

Actions

Answer Area

- Create a database master key.
- Create a column master key.
- Open the symmetric key.
- Create a certificate.
- Update Col1.
- Create a symmetric key.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

Reference:

<https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/>

NEW QUESTION 171

- (Exam Topic 5)

You have a 50-TB Microsoft SQL Server database named DB1.

You need to reduce the time it takes to perform database consistency checks of DB1.

Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area.

NOTE:Each correct selection is worth one point.

DBCC CHECKDB ([DB1],

▼
NOINDEX
REPAIR_FAST
REPAIR_REBUILD

 with

▼
ALL_ERRORMSGs
NO_INFOMSGs
PHYSICAL_ONLY

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated with low confidence

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-ser>

NEW QUESTION 173

- (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 175

- (Exam Topic 4)

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements. Which Azure Storage functionality should you include in the solution?

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

Answer: C

Explanation:

The lifecycle management policy lets you:

Delete blobs, blob versions, and blob snapshots at the end of their lifecycles Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

NEW QUESTION 178

- (Exam Topic 3)

Which windowing function should you use to perform the streaming aggregation of the sales data?

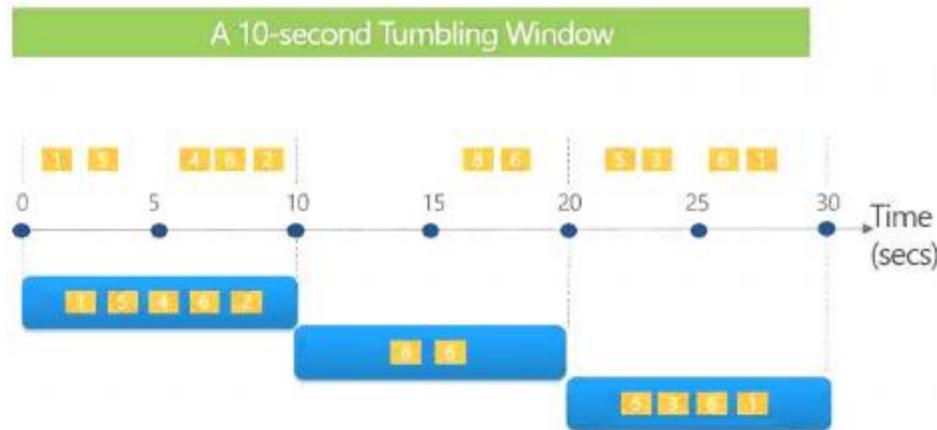
- A. Sliding
- B. Hopping
- C. Session
- D. Tumbling

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

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

NEW QUESTION 181

- (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.	<input type="radio"/>	<input type="radio"/>
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	<input type="radio"/>	<input type="radio"/>
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	<input type="radio"/>	<input type="radio"/>

- 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 186

- (Exam Topic 2)

Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

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

Answer: A

Explanation:

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

NEW QUESTION 188

- (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 192

- (Exam Topic 1)

You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Create and configure a schedule.
- Create a SQL Server Agent job.
- Publish the runbook.
- Create an Azure Automation account.
- Import the SqlServer module.
- Create a runbook that runs a PowerShell script.
- Run `sp_add_jobserver`.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Automating Azure SQL DB index and statistics maintenance using Azure Automation:

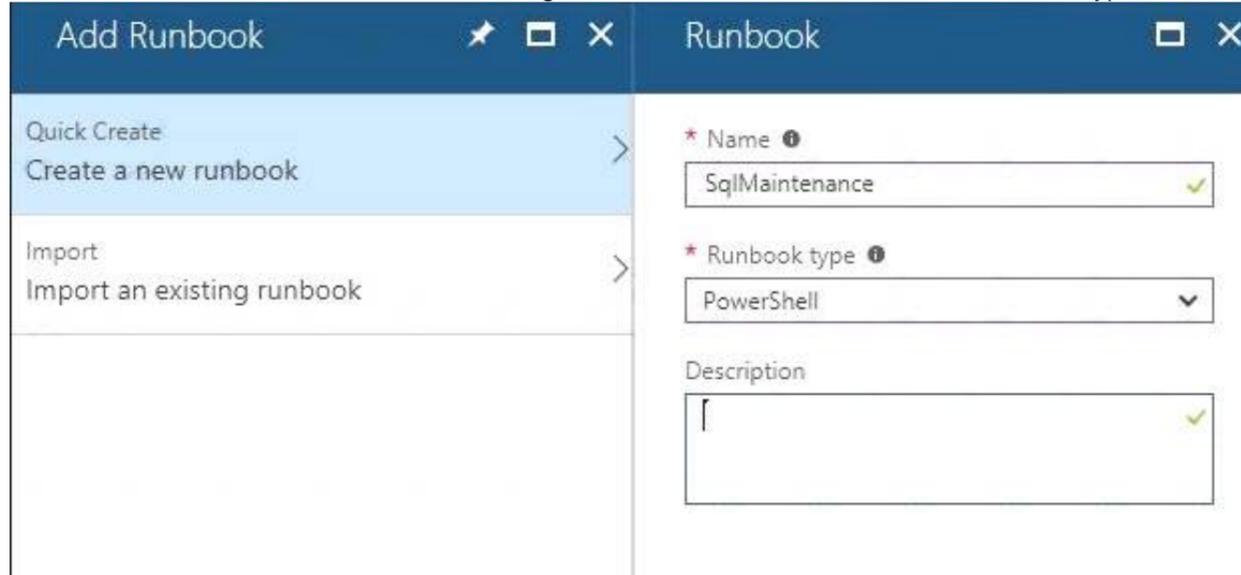
- * 1. Create Azure automation account (Step 1)
- * 2. Import SQLServer module (Step 2)
- * 3. Add Credentials to access SQL DB

This will use secure way to hold login name and password that will be used to access Azure SQL DB

* 4. Add a runbook to run the maintenance (Step 3)

Steps:1. Click on "runbooks" at the left panel and then click "add a runbook"

* 2. Choose "create a new runbook" and then give it a name and choose "Powershell" as the type of the runbook and then click on "create"



* 5. Schedule task (Step 4)

Steps:1. Click on Schedules2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.

Reference:

<https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statist>

NEW QUESTION 196

- (Exam Topic 1)

You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.

What should you run as part of the implementation?

- A. CREATE LOGINand theFROM WINDOWSclause
- B. CREATE USERand theFROM CERTIFICATEclause
- C. CREATE USERand theFROM LOGINclause
- D. CREATE USERand theASYMMETRIC KEYclause
- E. CREATE USERand theFROM EXTERNAL PROVIDERclause

Answer: E

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.) Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user: CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql>

NEW QUESTION 200

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

DP-300 Practice Exam Features:

- * DP-300 Questions and Answers Updated Frequently
- * DP-300 Practice Questions Verified by Expert Senior Certified Staff
- * DP-300 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DP-300 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DP-300 Practice Test Here](#)