Microsoft.DP-300.v2023-08-09.q117

Exam Code:	DP-300
Exam Name:	Administering Relational Databases on Microsoft Azure
Certification Provider:	Microsoft
Free Question Number:	117
Version:	v2023-08-09
# of views:	1306
# of Questions views:	1170
https://www.freepdfdumps.com/Microsoft.DP-300.v2023-08-09.q117.html	

NEW QUESTION: 1
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. email notifications
B. log to table
C. output files
D. schedules
Answer: B (LEAVE A REPLY)

NEW QUESTION: 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 (LEAVE A REPLY)
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: 3
You configure version control for an Azure Data Factory instance as shown in the following exhibit.

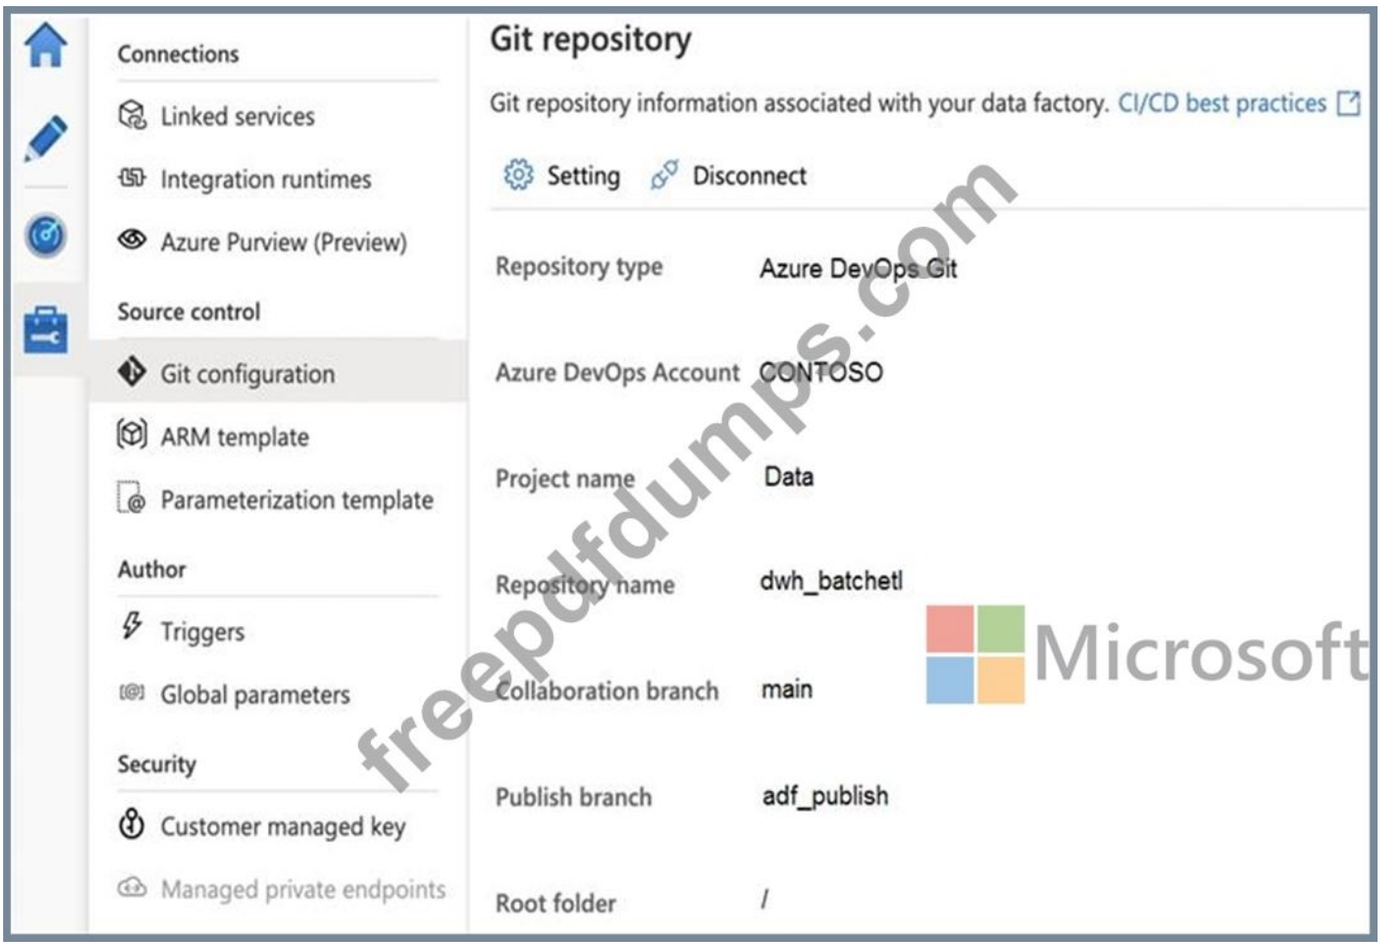

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Azure Resource Manager (ARM) templates or the pipeline assets as stored in

- /
- adf_publish
- main
- Parameterization template

A Data Factory Azure Resource Manager (ARM) template named contososales can be found in

- /contososales
- /dwh_batchetl/adf_publish/contososales
- /main

Answer:

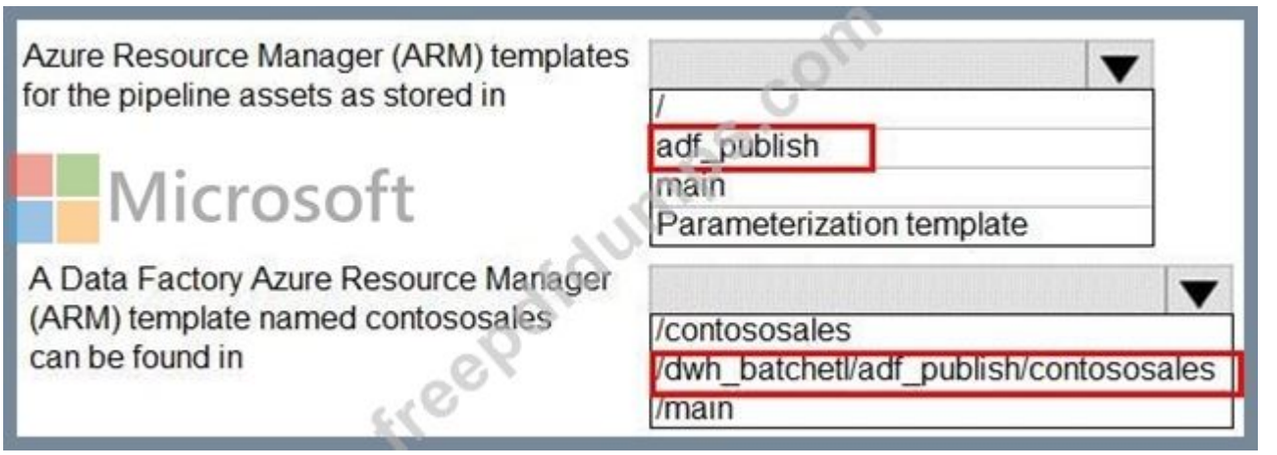

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/source-control

NEW QUESTION: 4
You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information.
You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers.
The solution must prevent all the salespeople from viewing or inferring the credit card information.
What should you include in the recommendation?
A. row-level security
B. data masking
C. Always Encrypted
D. column-level security
Answer: (SHOW ANSWER)
Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.
The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card.
Example:
XXXX-XXXX-XXXX-1234

NEW QUESTION: 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. private endpoints

B. a site-to-site VPN

C. service endpoints

D. virtual network peering

Answer: A (LEAVE A REPLY)

NEW QUESTION: 6

You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database.

A recent failure causes the master database to be lost.

You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

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

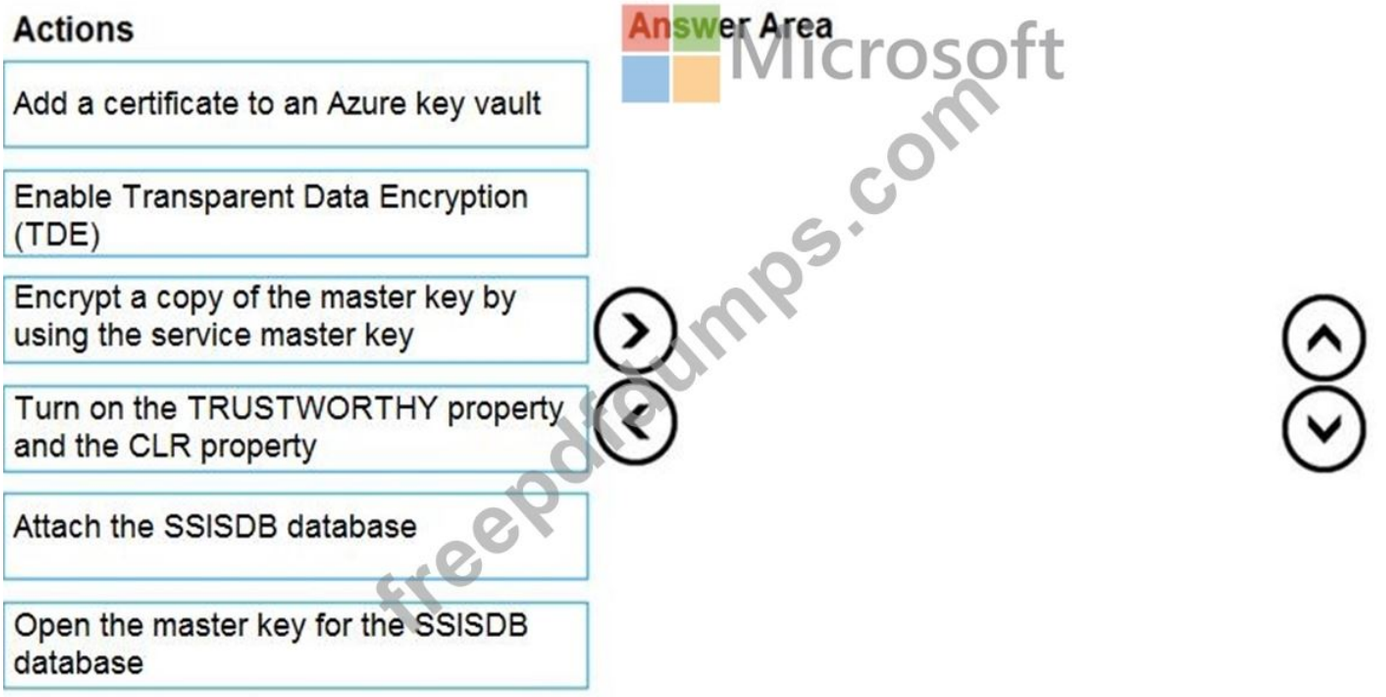

Answer:

Answer Area

Attach the SSISDB database
Turn on the TRUSTWORTHY property and the CLR property
Open the master key for the SSISDB database
Encrypt a copy of the mater key by using the service master key

1 - Attach the SSISDB database
2 - Turn on the TRUSTWORTHY property and the CLR property
3 - Open the master key for the SSISDB database
4 - Encrypt a copy of the mater key by using the service master key
Reference:
https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog

NEW QUESTION: 7
You have SQL Server on an Azure virtual machine named SQL1.
SQL1 has an agent job to back up all databases.
You add a user named dbadmin1 as a SQL Server Agent operator.
You need to ensure that dbadmin1 receives an email alert if a job fails.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

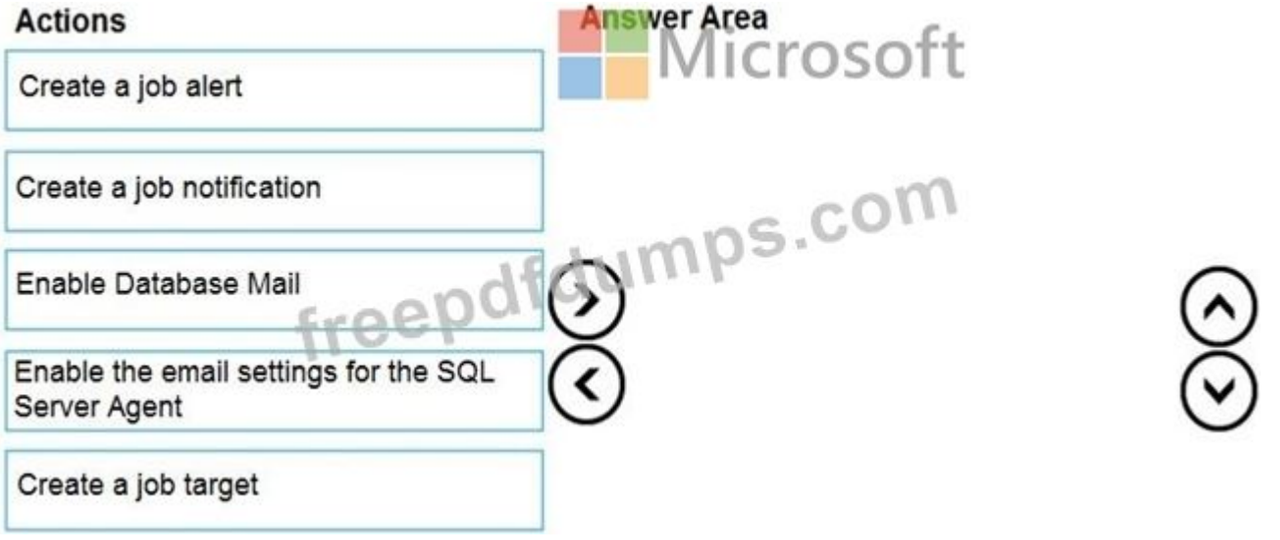

Answer:

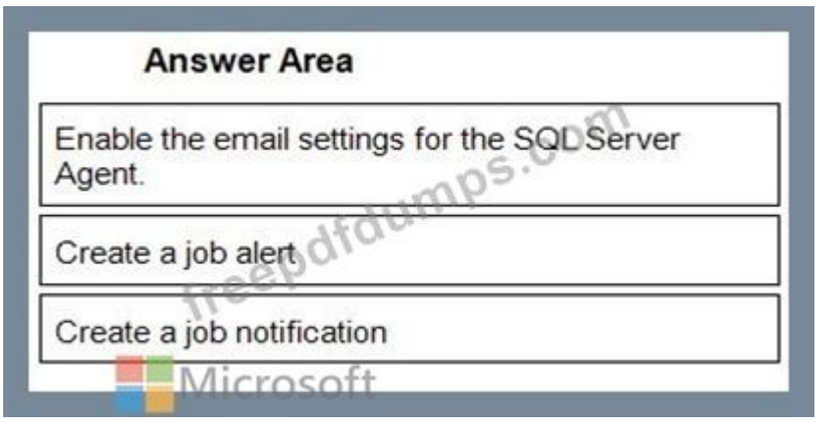

1 - Enable the email settings for the SQL Server Agent.
2 - Create a job alert
3 - Create a job notification
Reference:
https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status
https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator

NEW QUESTION: 8
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 (LEAVE A REPLY)
Azure Databricks makes a distinction between all-purpose clusters and job clusters. You use all-purpose clusters to analyze data collaboratively using interactive notebooks. You use job clusters to run fast and robust automated jobs.
The Azure Databricks job scheduler creates a job cluster when you run a job on a new job cluster and terminates the cluster when the job is complete.
Reference:
https://docs.microsoft.com/en-us/azure/databricks/clusters

NEW QUESTION: 9
You have two Azure virtual machines named VM1 and VM2 that run Windows Server 2019. VM1 and VM2 each host a default Microsoft SQL Server 2019 instance. VM1 contains a database named DB1 that is backed up to a file named D:\DB1.bak.
You plan to deploy an Always On availability group that will have the following configurations:
VM1 will host the primary replica of DB1.
VM2 will host a secondary replica of DB1.
You need to prepare the secondary database on VM2 for the availability group.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

Answer:

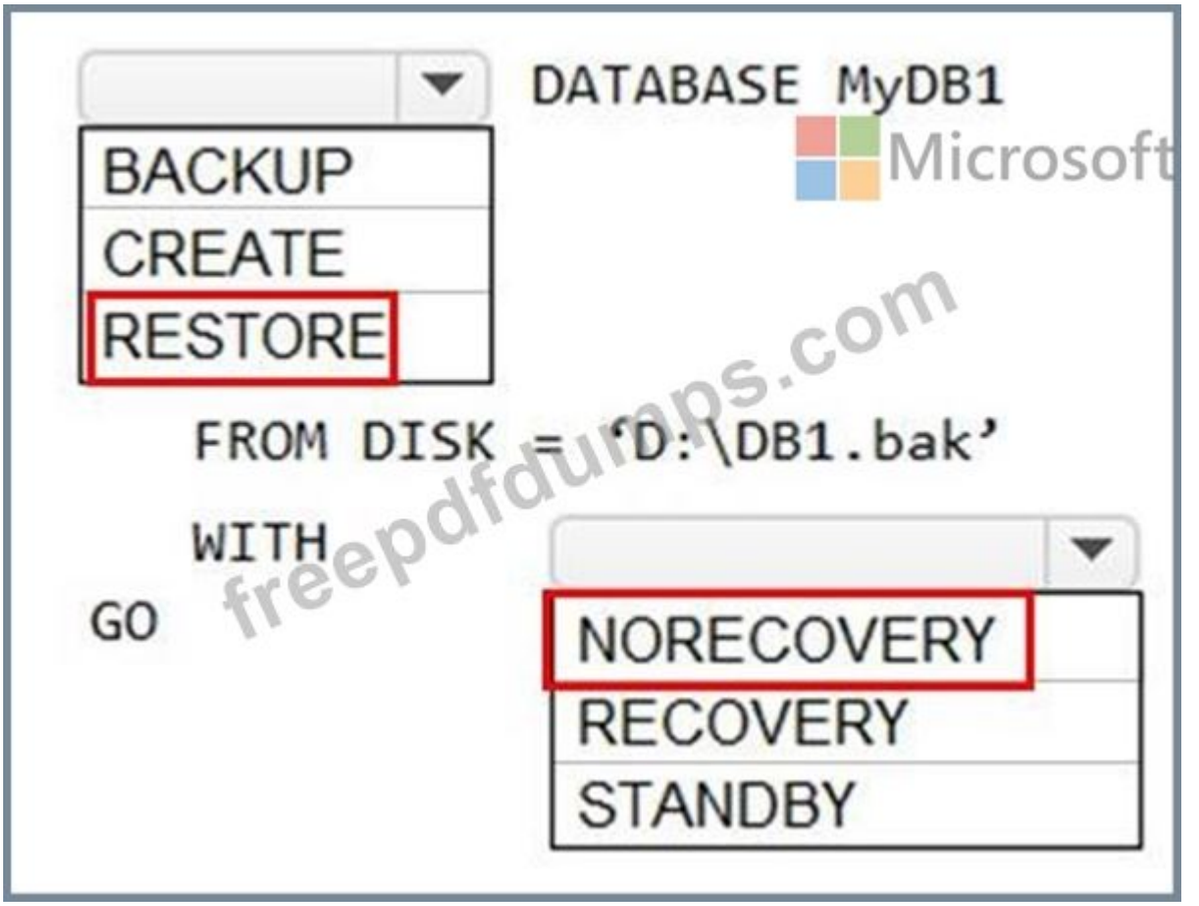

Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server?view=sql-server-ver15

NEW QUESTION: 10

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. Column2
B. Column3
C. Column1
D. Column4
E. Column5
Answer: (SHOW ANSWER)

NEW QUESTION: 11
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
SQL1	SQL Server on Azure Virtual Machines	**Not applicable**
db1	Microsoft SQL Server database	Hosted on SQL1
mysqlbackups	General purpose v2 storage account	**Not applicable**

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.
Which three Transact-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

```
RESTORE DATABASE db2 FROM URL = URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup', RECOVERY,
     MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
     MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

```
RESTORE DATABASE db2 FROM URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH RECOVERY,
     MOVE 'db1_mdf' TO
'D:\Data\db2_mdf.mdf',
     MOVE 'db1_log' TO
'D:\Logs\db2_log.ldf'
```

```
CREATE CREDENTIAL
[https://mysqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

```
BACKUP DATABASE db1
TO URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_2.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://mysqlbackups.blob.core.windows.net
```

Answer Area

```
/backups/db1_4.bak'

GO
```

```
CREATE CREDENTIAL [sqlbackup] WITH IDENTITY
=
'sqlsamplebackup'
,SECRET = '<mystorageaccountaccesskey>';
GO
```

Answer:

Answer Area

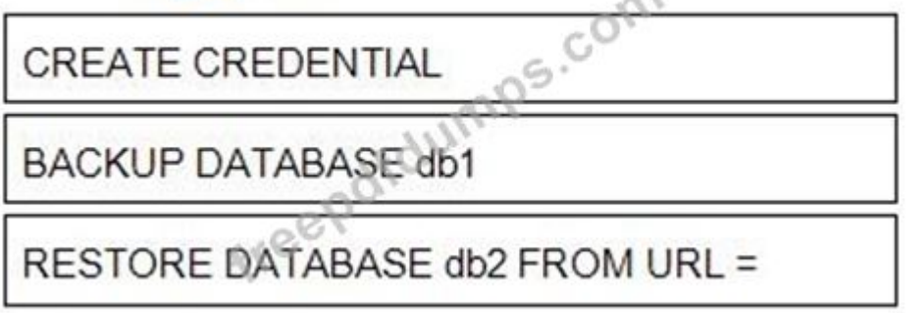

1 - CREATE CREDENTIAL
2 - BACKUP DATABASE db1
3 - RESTORE DATABASE db2 FROM URL =
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-server-ver15

NEW QUESTION: 12
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 dat a. 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 (LEAVE A REPLY)
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: 13
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 (LEAVE A REPLY)
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: 14
You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1.
Several users execute adhoc queries to DW1 concurrently.
You regularly perform automated data loads to DW1.
You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run.
What should you do?
A. Assign a smaller resource class to the automated data load queries.
B. Create sampled statistics to every column in each table of DW1.
C. Assign a larger resource class to the automated data load queries.
D. Hash distribute the large fact tables in DW1 before performing the automated data loads.
Answer: C (LEAVE A REPLY)
The performance capacity of a query is determined by the user's resource class.
Smaller resource classes reduce the maximum memory per query, but increase concurrency.
Larger resource classes increase the maximum memory per query, but reduce concurrency.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-

for-workloadmanagement

NEW QUESTION: 15

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: B,C (LEAVE A REPLY)

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell#best-practices-for-sql-database

https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azure-portal

NEW QUESTION: 16

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 (LEAVE A REPLY)

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: 17
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. an Azure SQL Managed Instance link
B. availability groups for SQL Server on Azure Virtual Machines
C. transactional replication to an Azure SQL Managed Instance
D. database mirroring on an instance of SQL Server on Azure Virtual Machines
Answer: C (LEAVE A REPLY)

NEW QUESTION: 18
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.

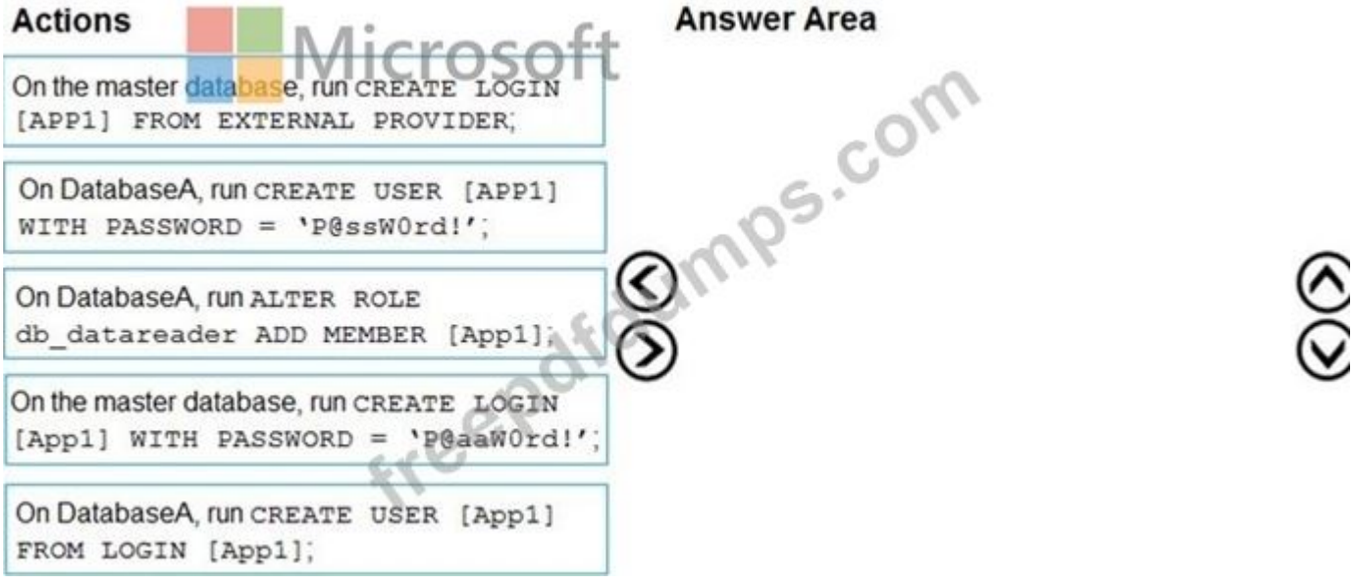

Answer:

1 - On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'.
2 - On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]
3 - On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]
Reference:
https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/

NEW QUESTION: 19
You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.
You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.
You need to ensure that webapp1 can identity the changes to Table'. The solution must meet the following requirements:
* Minimize compute times.
* Minimize storage.
Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

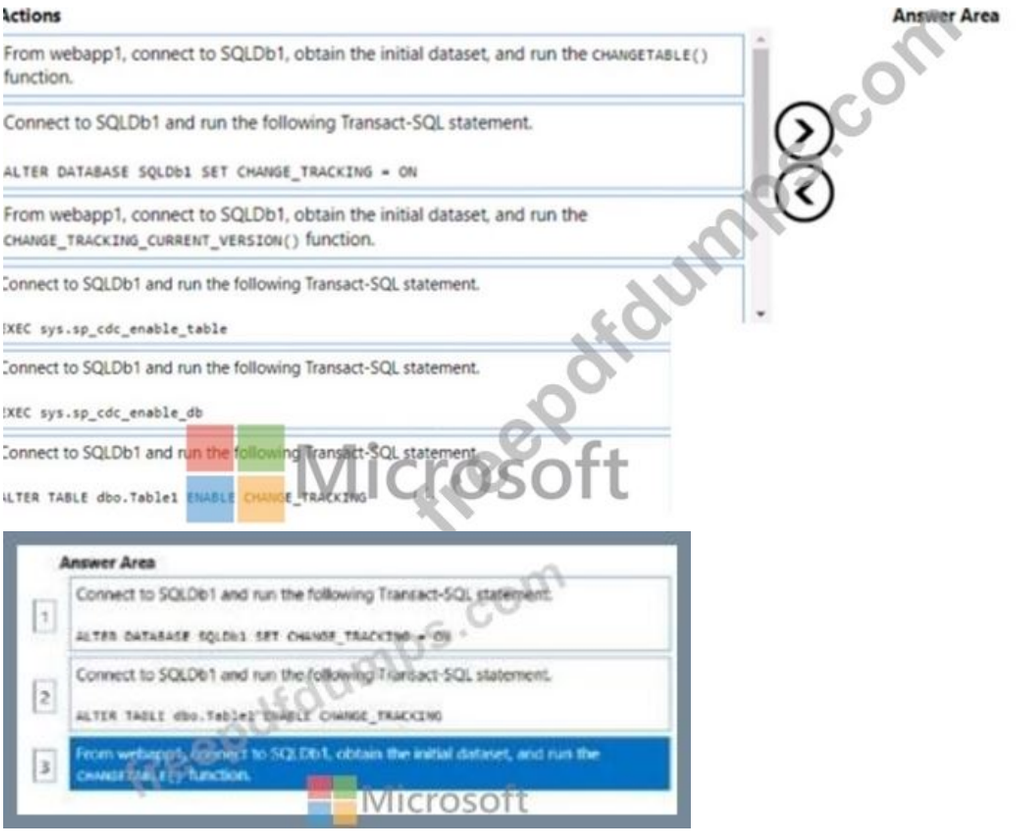

Answer:

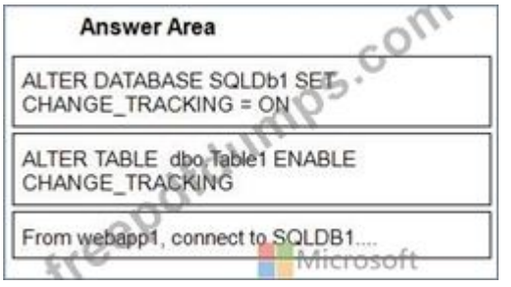

1 - ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON

2 - ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING

3 - From webapp1, connect to SQLDB1....

NEW QUESTION: 20

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?

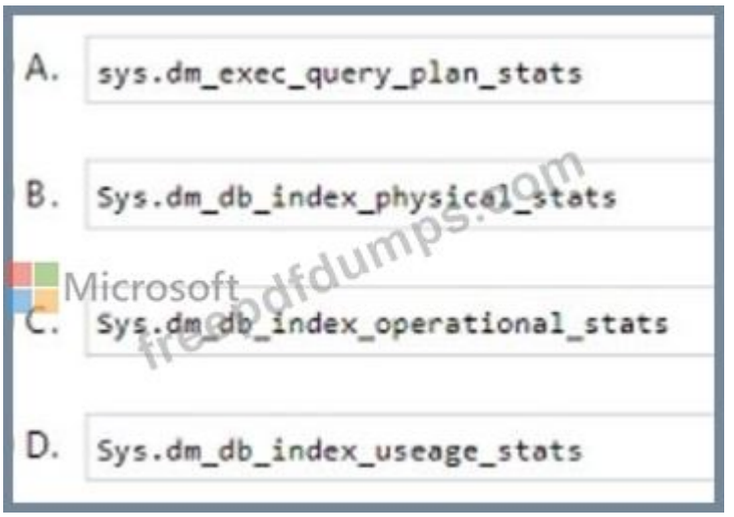

A. Option B
B. Option C
C. Option D
D. Option A
Answer: (SHOW ANSWER)

NEW QUESTION: 21
You have an Azure SQL managed instance named MI1.
You need to implement automatic tuning for the databases of MI1.
What should you do?
A. From the Azure portal, configure automatic tuning.
B. Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property.
C. Use Transact-SQL to enable the force_last_good_plan option.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 22
You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in fiveminute intervals and report only events that arrive during the interval.
The output will be sent to a Delta Lake table.
Which output mode should you use?
A. complete
B. append
C. update
Answer: (SHOW ANSWER)
Complete mode: You can use Structured Streaming to replace the entire table with every batch.
Incorrect Answers:
B: By default, streams run in append mode, which adds new records to the table.
Reference:
https://docs.databricks.com/delta/delta-streaming.html

NEW QUESTION: 23

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.

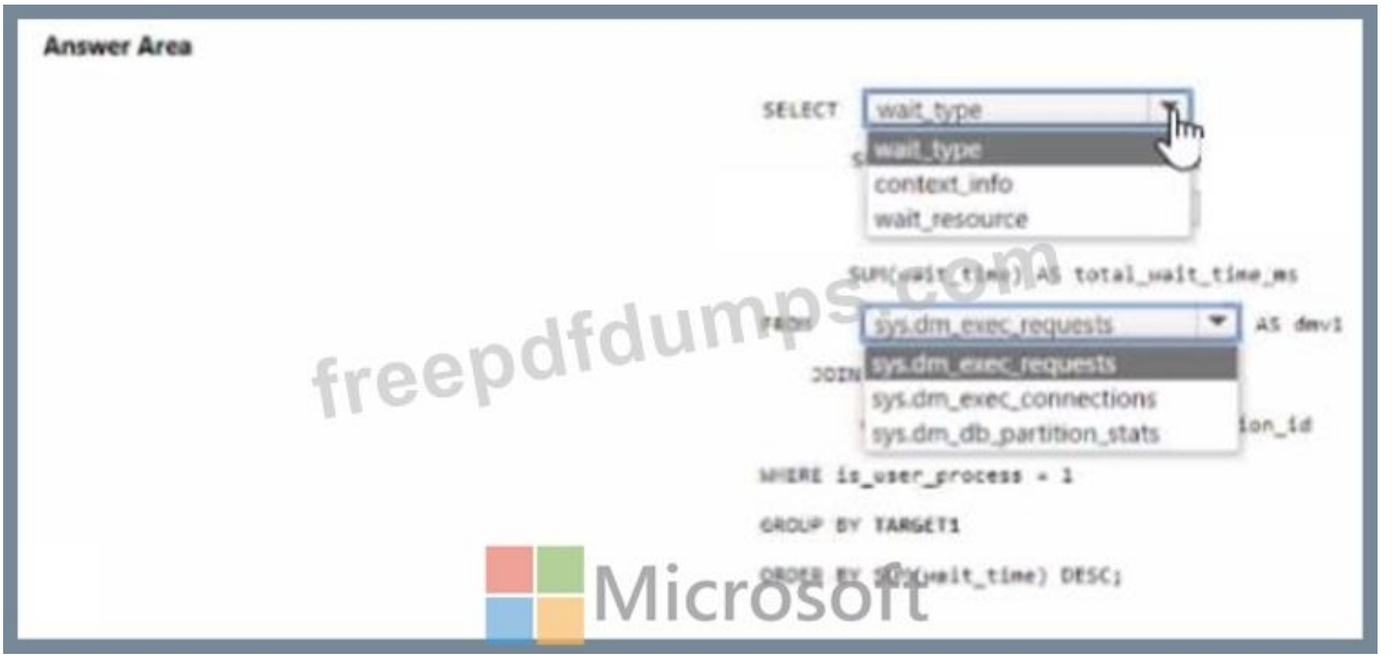

Answer:

NEW QUESTION: 24

You have a new Azure SQL database. The database contains a column that stores confidential information.

You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.

NOTE: Each correct selection is worth one point.

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

A. Turn on auditing and write audit logs to an Azure Storage account.

B. Add extended properties to the column.
C. Turn on Advanced Data Security for the Azure SQL server.
D. Apply sensitivity labels named Highly Confidential to the column.
E. Turn on Azure Advanced Threat Protection (ATP).
Answer: (SHOW ANSWER)
C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities. ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.
A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data_sensitivity_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview

NEW QUESTION: 25
You need to design an analytical storage solution for the transactional dat a. The solution must meet the sales transaction dataset requirements.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

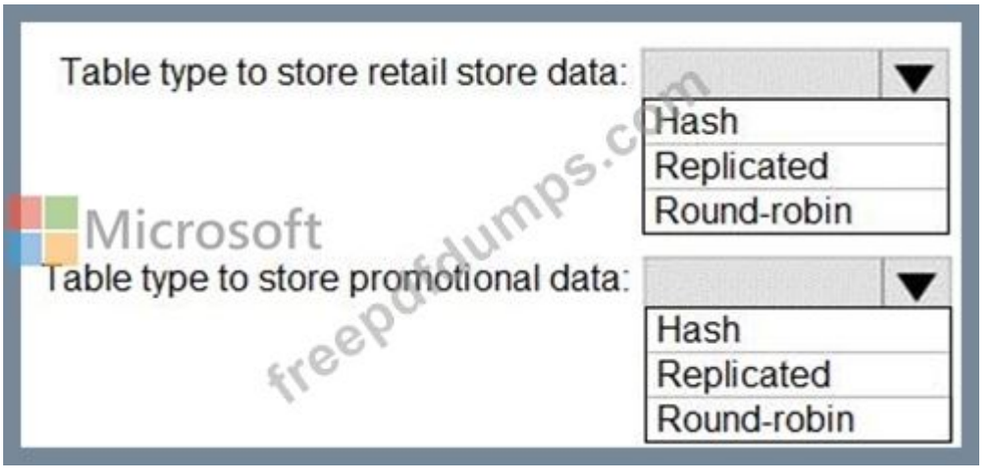

Answer:

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

NEW QUESTION: 26
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.
You scale up the virtual machine to 8 vCPUSs and 64 GB of memory.
You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?
A. 2
B. 4
C. 8
D. 64
Answer: C (LEAVE A REPLY)
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

NEW QUESTION: 27
You have an Azure SQL database named db1 that contains an Azure Active Directory (Azure AD) user named user1.
You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context.
How should you complete the Transact-SQL statement? To answer, select the appropriate

options in the answer area.

NOTE: Each correct selection is worth one point.

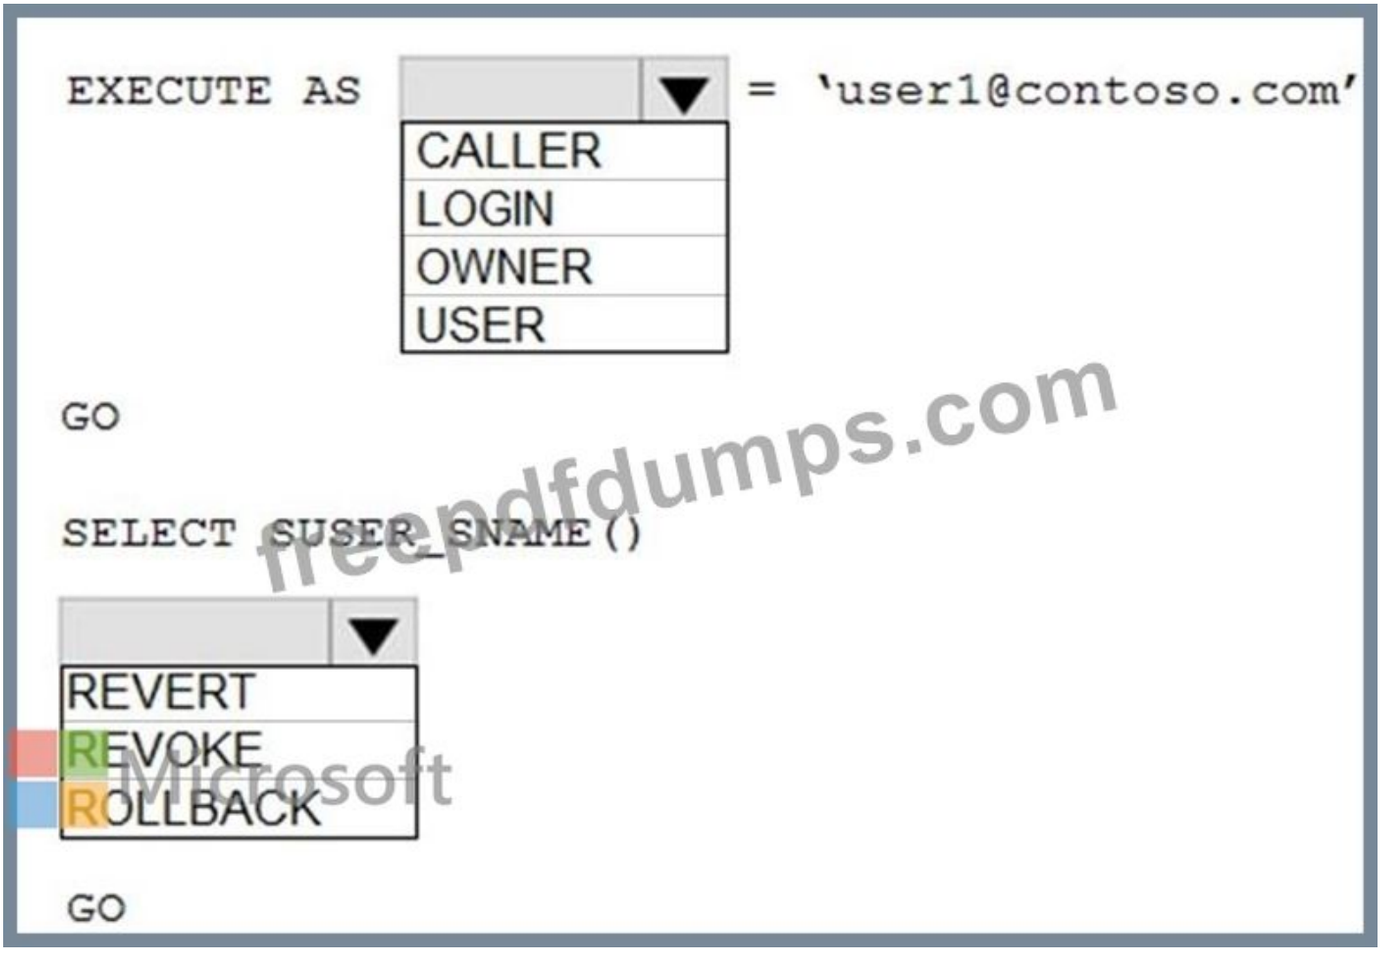

Answer:

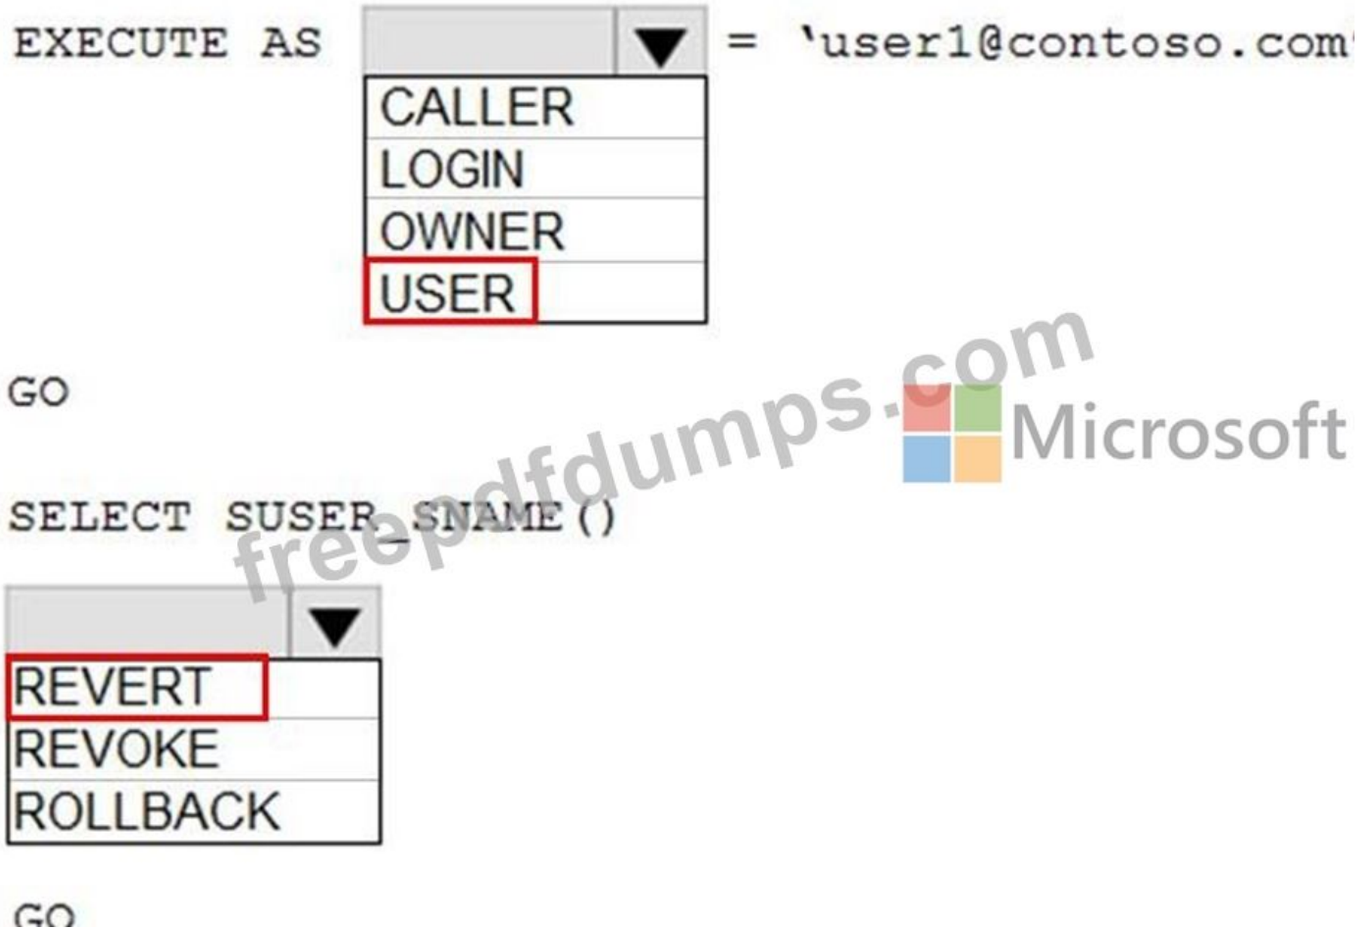

Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15
https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15

NEW QUESTION: 28
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: (SHOW ANSWER)
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)
Access control lists (ACL) Data Lake Storage Gen2 supports the following authorization mechanisms:
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: 29
You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.
You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:
The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
Administrative effort must be minimized.
What should you use to host the databases?
A. SQL Server on Azure Virtual Machines
B. an Azure SQL Database elastic pool
C. a single Azure SQL database
D. Azure SQL Managed Instance
Answer: D (LEAVE A REPLY)

NEW QUESTION: 30
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: (SHOW ANSWER)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database
NEW QUESTION: 31

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 (LEAVE A REPLY)
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: 32
You have the following Transact-SQL query.

```
    [file_id] AS [File ID],
    [type] AS [File Type],
    substring([physical_name], 1,1) AS [Drive],
    [name] AS [Logical Name],
    [physical_name] AS [Physical Name],
    CAST([size] as DECIMAL(38,0))/128.0 AS [ColumnA],
    CAST(FILEPROPERTY([name], 'SpaceUsed') AS DECIMAL(38,0))/128.0 AS
ColumnB],
    (CAST([size] AS DECIMAL(38,0))/128.0) - (CAST(FILEPROPERTY([name]
SpaceUsed') AS DECIMAL (38,0))/128.0) AS [ColumnC],
    [max_size] AS [ColumnD],
    [is_percent_growth] AS [Percent Growth Enabled],
    [growth] AS [Growth Rate],
    SYSDATETIME() AS [Current Date]
FROM sys.database_files;
```

Which column returned by the query represents the free space in each file?
A. ColumnA
B. ColumnB
C. ColumnC
D. ColumnD
Answer: C (LEAVE A REPLY)
Example:
Free space for the file in the below query result set will be returned by the FreeSpaceMB column.
SELECT DB_NAME() AS DbName,

name AS FileName,
type_desc,
size/128.0 AS CurrentSizeMB,
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB
FROM sys.database_files WHERE type IN (0,1); Reference:
https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/

NEW QUESTION: 33
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: A,E (LEAVE A REPLY)
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.
Incorrect Answers:
C: individual SQL Managed Instance databases cannot be added to or removed from a failover group.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

NEW QUESTION: 34
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

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

NOTE: Each correct selection is worth one point.

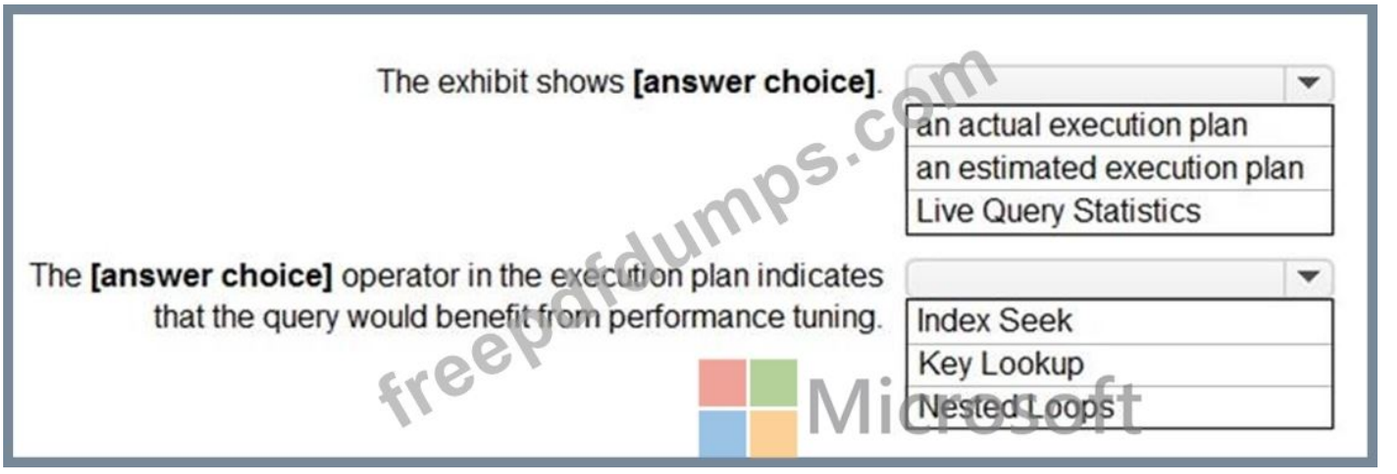

Answer:

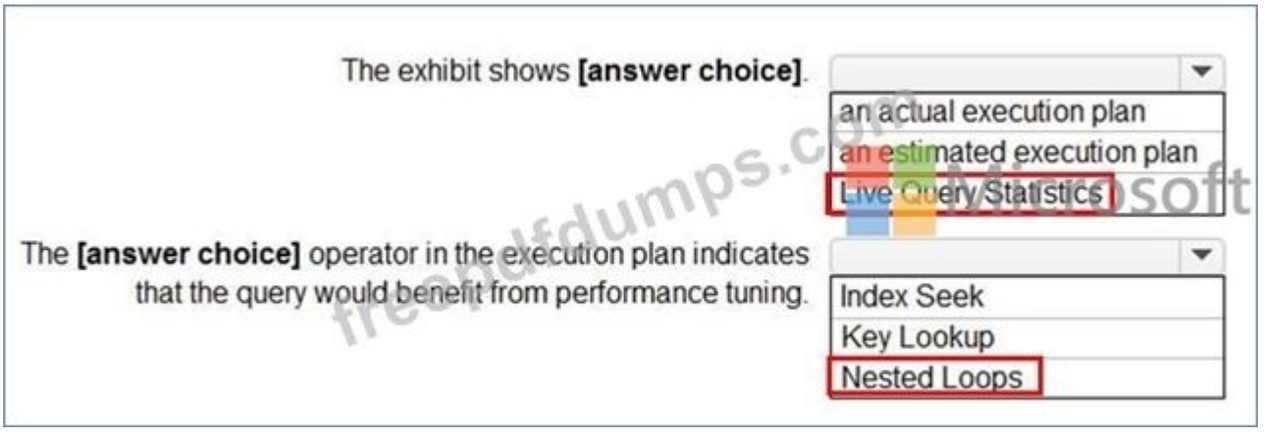

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15

NEW QUESTION: 35

You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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

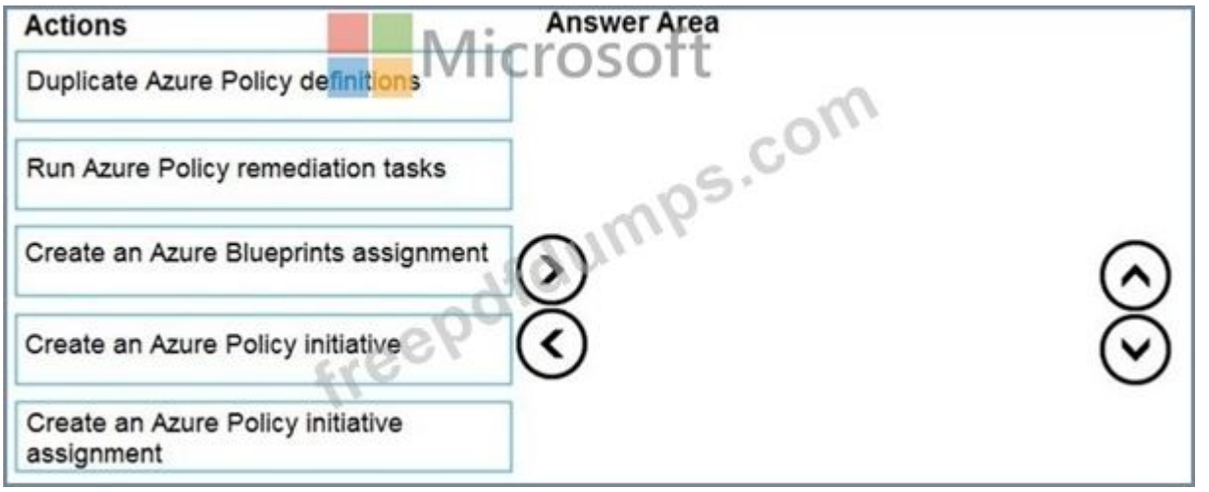

Answer:

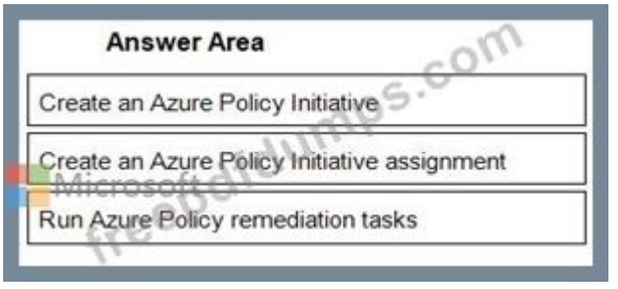

1 - Create an Azure Policy Initiative

2 - Create an Azure Policy Initiative assignment

3 - Run Azure Policy remediation tasks

Reference:

https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

NEW QUESTION: 36

You have an Azure Data Factory pipeline that is triggered hourly.

The pipeline has had 100% success for the past seven days.

The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

```
ErrorCode=UserErrorFileNotFound,
'Type=Microsoft.DataTransfer.Common.Shared.HybridDeliveryException,Message=ADLS
Gen2 operation failed for: Operation returned an invalid status code
'NotFound'. Account: 'contosoproduksouth' FileSystem: wwi.Path:
'BIKES/CARBON/year=2021/month=01/day=10/hour=06'. ErrorCode:
'PathNotFound'.Message: 'The specified path does not exist.'. RequestId:
'6d269b78-901f-001b-4924-e7a7bc000000'. TimeStamp: 'Sun, 10 Jan 2021 07:45:05
```

What is a possible cause of the error?

A. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.

B. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.

C. The pipeline was triggered too early.

D. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 37

You have an Azure subscription.

You need to deploy an Azure SQL managed instance by using an Azure Resource Manager (ARM) template. The solution must meet the following requirements:

The SQL managed instance must be assigned a unique identity.

The SQL managed instance must be available in the event of an Azure datacenter outage.

How should you complete the template? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

NEW QUESTION: 38

Your company uses Azure Stream Analytics to monitor devices.
The company plans to double the number of devices that are monitored.
You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.
Which metric should you monitor?
A. Input Deserialization Errors
B. Late Input Events
C. Early Input Events
D. Watermark delay
Answer: D (LEAVE A REPLY)
The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.
The watermark delay metric can rise due to:
1. Not enough processing resources in Stream Analytics to handle the volume of input events.
2. Not enough throughput within the input event brokers, so they are throttled.
3. Output sinks are not provisioned with enough capacity, so they are throttled.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling

NEW QUESTION: 39
You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 contains a table named CustomerPII.
You need to record whenever users query the CustomerPII table.
Which two options should you enable? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. server audit specification
B. SQL Server audit
C. database audit specification
D. a server principal
Answer: (SHOW ANSWER)
An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):
A server policy applies to all existing and newly created databases on the server.
If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings.
Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.
Note:
A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.
The Server Audit Specification object belongs to an audit.
Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/auditing-overview

NEW QUESTION: 40

You have an Azure SQL database.

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

You run the select * from sys.database_scoped_configurations Transact-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: (SHOW ANSWER)

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.

Incorrect Answers:

A: LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and earlier version independent of the compatibility level of the database. The default is OFF, which sets the query optimizer cardinality estimation model based on the compatibility level of the database.

B: QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY }

Enables or disables query optimization hotfixes regardless of the compatibility level of the database. The default is OFF, which disables query optimization hotfixes that were released after the highest available compatibility level was introduced for a specific version (post-RTM).

Reference:

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

NEW QUESTION: 41

You have an on-premises Microsoft SQL server that uses the FileTables and Filestream features.

You plan to migrate to Azure SQL.
Which service should you use?
A. Azure SQL Database
B. SQL Server on an Azure Virtual Machine
C. Azure SQL Managed Instance
D. Azure Database for MySQL
Answer: B (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database-overview

NEW QUESTION: 42
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.
What should you use?
A. An Azure SQL Database elastic pool
B. SQL Server on Azure virtual machines
C. an Azure SQL Database managed instance
D. Azure SQL databases
Answer: (SHOW ANSWER)
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: 43
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 (LEAVE A REPLY)
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: 44

You need to use an Azure Resource Manager ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files How should you complete the template? To answer, select the appropriate options in the answer are a.

NOTE: Each correct selection is worth one point.

Answer:

NEW QUESTION: 45

Note: This question is part of a series of questions that present the same scenario. Each question

in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:
* During normal operations, provide at least two readable copies of Sales.
* Ensure that Sales remains available if a datacenter fails.
Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication.
Does this meet the goal?
A. Yes
B. No
Answer: B (LEAVE A REPLY)
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: 46
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 (LEAVE A REPLY)

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: 47

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.	○	○

Answer:

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: 48

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. Always On Availability Group

B. distributed availability groups

C. Azure Database Migration Service

D. database mirroring

Answer: (SHOW ANSWER)

NEW QUESTION: 49

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1.

You need to enable Resource Governor for a App1. The solution must meet the following requirements:

App1 must be able to consume all available CPU resources.

App1 must have at least half of the available CPU resources always available.

Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

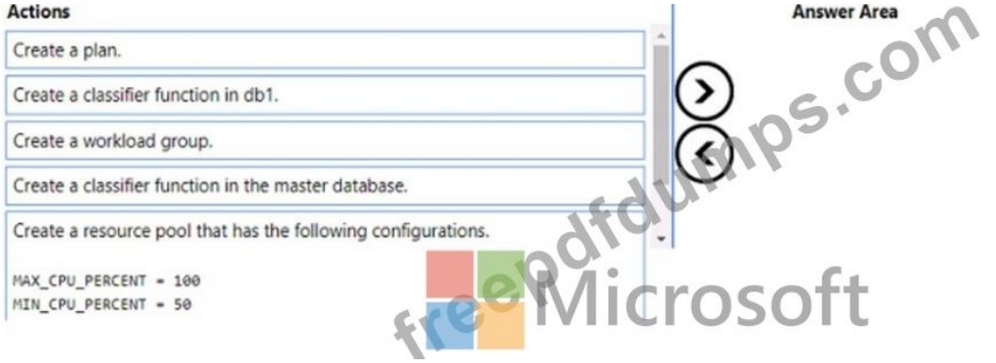

Answer:

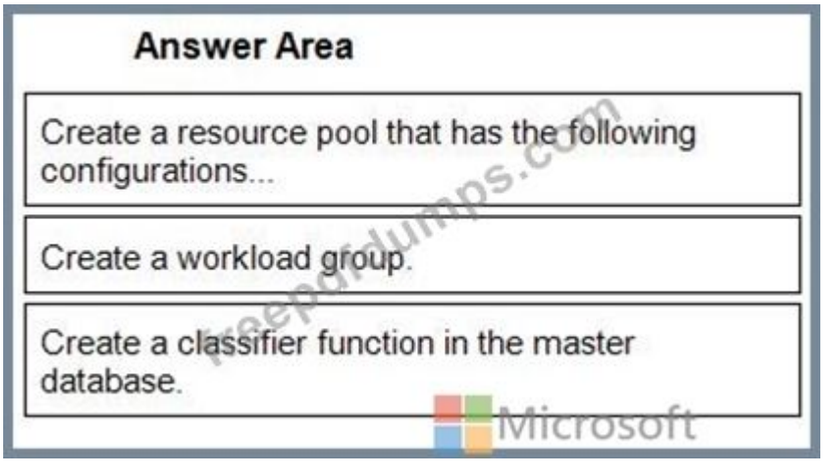

1 - Create a resource pool that has the following configurations...

2 - Create a workload group.

3 - Create a classifier function in the master database.

NEW QUESTION: 50

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.
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.	○	○
An alert notification was sent after the failure of Activity3 in PipelineA.	○	○
An alert notification was sent after the failure of Activity1 in PipelineB.	○	○

Answer:

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	○	○
An alert notification was sent after the failure of Activity3 in PipelineA.	○	○
An alert notification was sent after the failure of Activity1 in PipelineB.	○	○

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

NEW QUESTION: 51
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: B,D (LEAVE A REPLY)
Reference:
https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-ina-data-warehouse

NEW QUESTION: 52
You have An Azure SQL managed instance.
You need to configure the SQL Server Agent service to email job notifications.
Which statement should you execute?
A)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
```

B)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'application_dbmail_profile';
```

C)

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
```

```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sys_dbmail_profile';
```

A. Option D
B. Option B
C. Option A
D. Option C
Answer: B (LEAVE A REPLY)

NEW QUESTION: 53
You have SQL Server on an Azure virtual machine.
You review the query plan shown in the following exhibit.

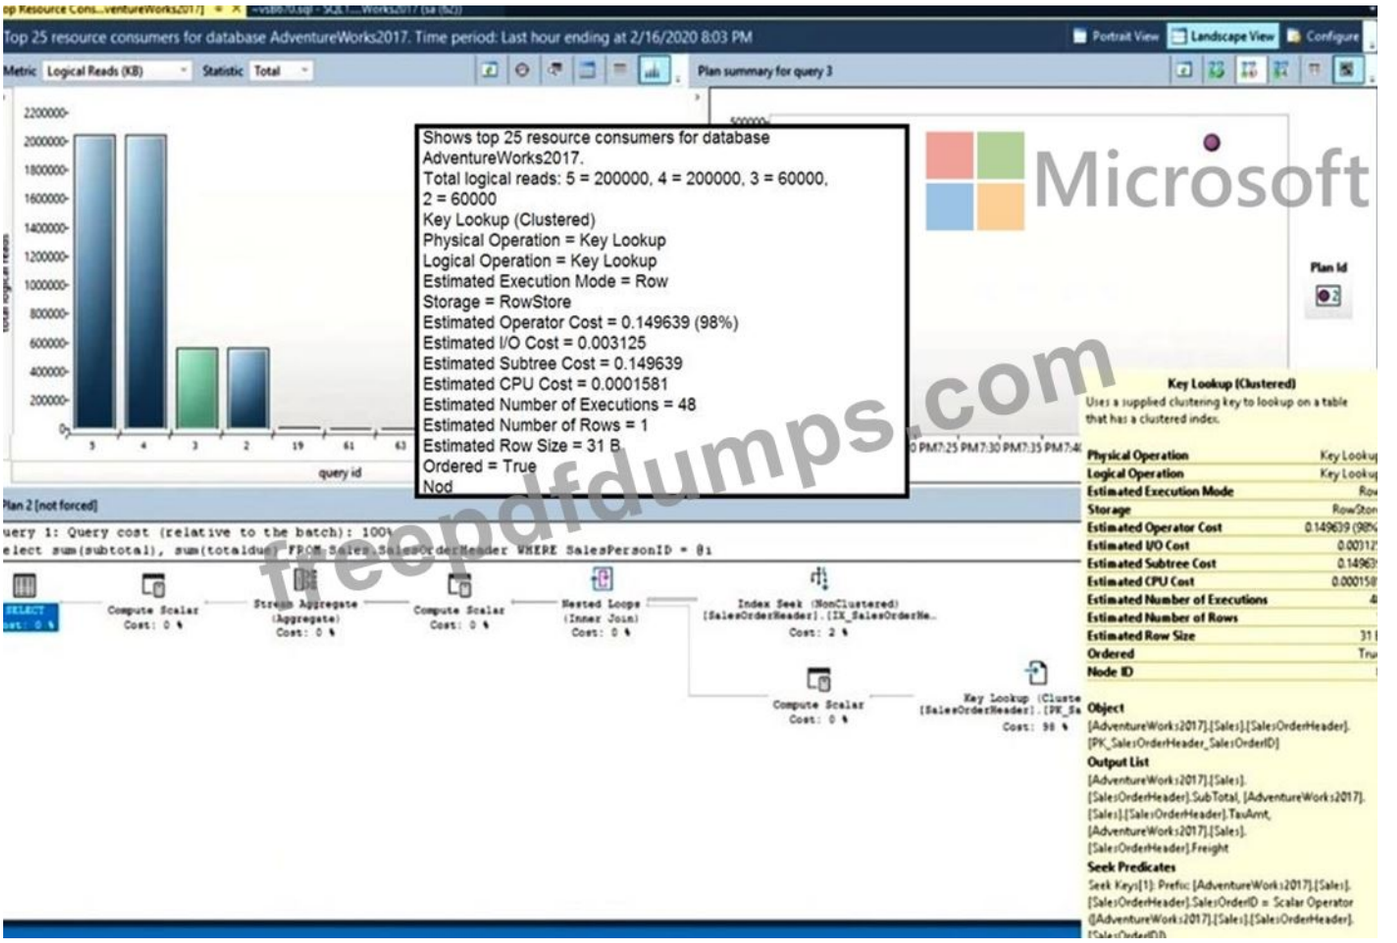

For each of the following statements, select yes if the statement is true. Otherwise, select no.

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.	○	○

Answer:

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.	○	○

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store

NEW QUESTION: 54
You create five Azure SQL Database instances on the same logical server.
In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.
User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.
You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.
What should you do?
A. Create User1 in the master database.
B. Assign User1 the db_datareader role for the master database.
C. Assign User1 the db_datareader role for the databases that Userl creates.
D. Grant select on sys.databases to public in the master database.
Answer: A (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage

NEW QUESTION: 55
You have an instance of SQL Server on Azure Virtual Machines.
You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The solution must use the principle of least privilege.
Which role should you assign to User1?
A. sysadmin
B. SQLAgentReaderRole
C. SQLAgentUserRole
D. SQLAgentOperatorRole
Answer: C (LEAVE A REPLY)

NEW QUESTION: 56
You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.
Which resource provider should you enable?
A. Microsoft.EventHub
B. Microsoft.EventGrid
C. Microsoft.Sql
D. Microsoft.Automation
Answer: (SHOW ANSWER)
Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data

Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger

NEW QUESTION: 57
You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.
You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

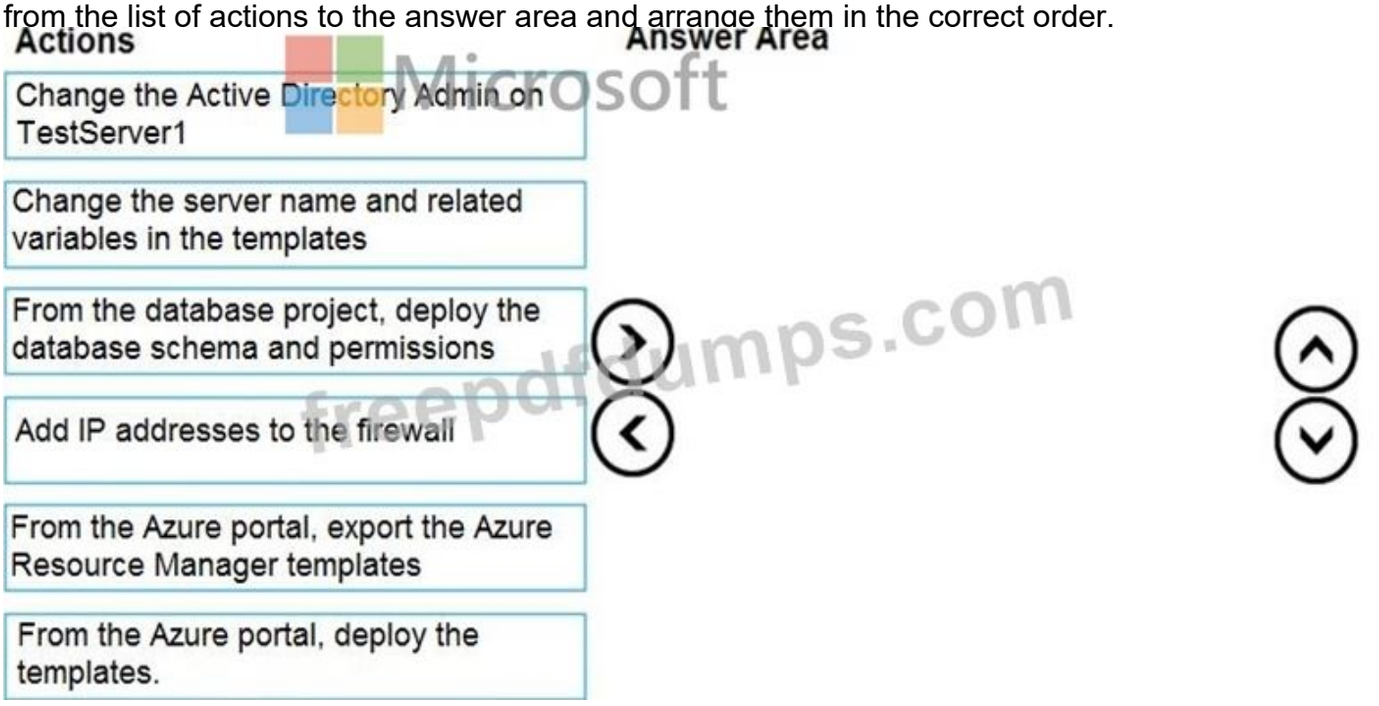

Answer:

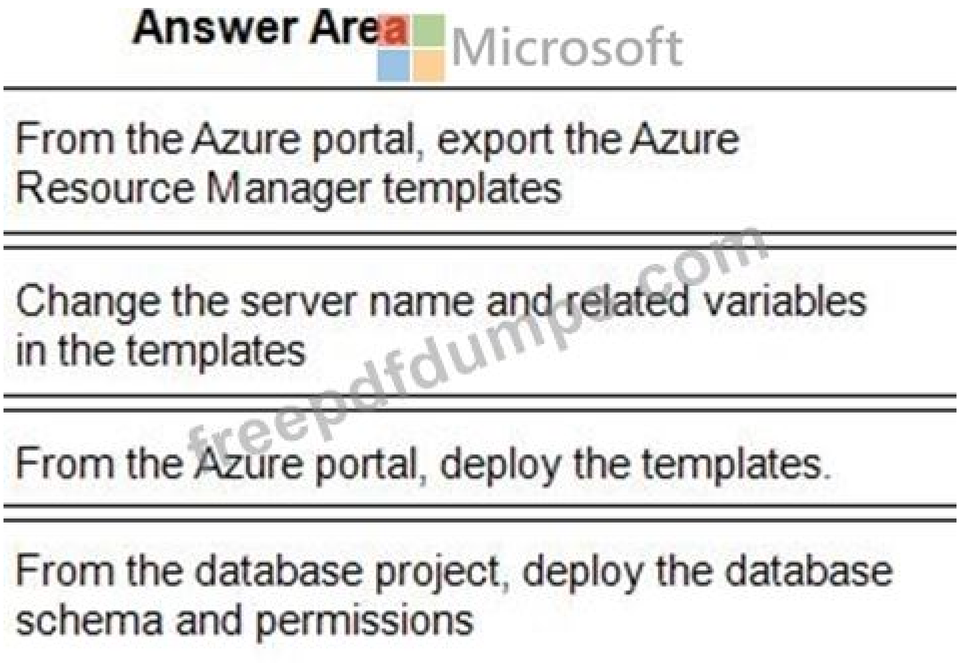

1 - From the Azure portal, export the Azure Resource Manager templates
2 - Change the server name and related variables in the templates
3 - From the Azure portal, deploy the templates.
4 - From the database project, deploy the database schema and permissions

NEW QUESTION: 58

You have an Azure subscription that contains an Azure SQL database.
The database fails to respond to queries in a timely manner.
You need to identify whether the issue relates to resource_semaphore waits.
How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

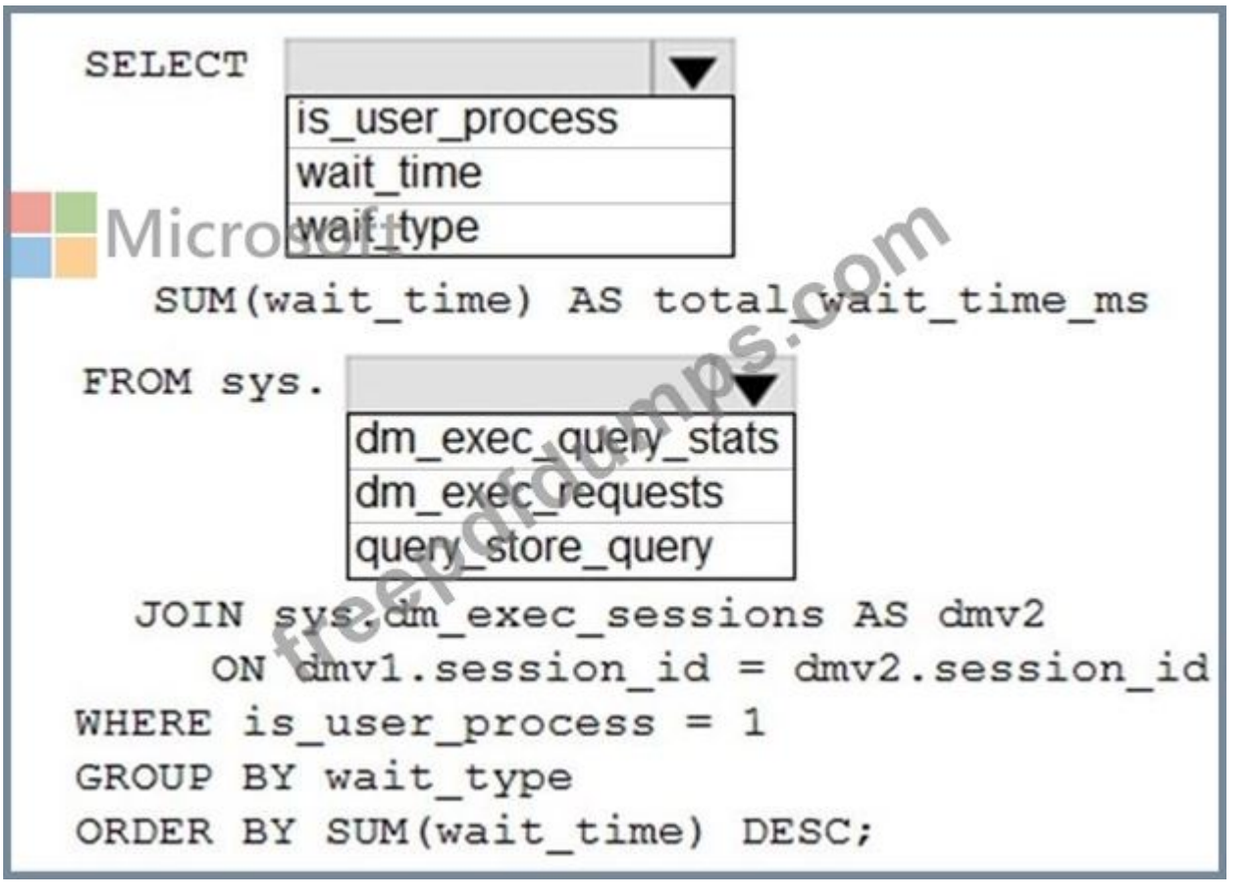

Answer:

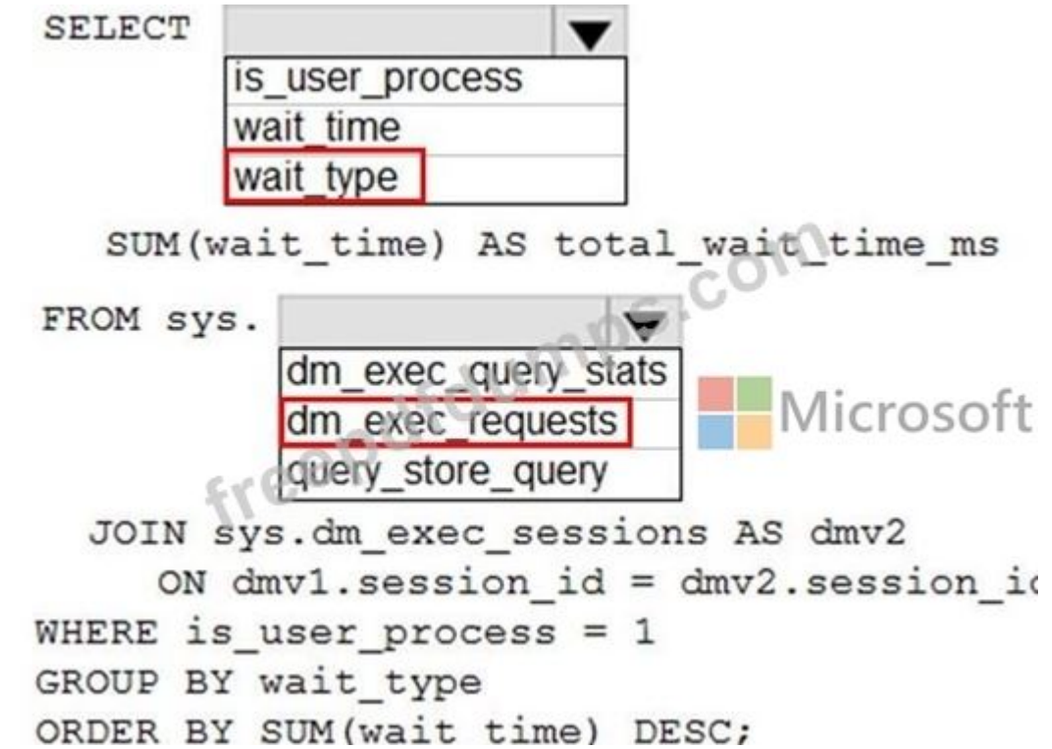

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs

NEW QUESTION: 59

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state.

What should you do?

A. Double the value of Data Flush interval

B. Decrease by half the value of Data Flush Interval

C. Double the value of Statistics Collection Interval

D. Decrease by half the value of Statistics Collection interval

Answer: (SHOW ANSWER)

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the maximum size limit between storage size checks, it transitions to read-only mode.

Incorrect Answers:

C: Statistics Collection Interval: Defines the level of granularity for the collected runtime statistic, expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues. Keep in mind that the value directly affects the size of Query Store data.

Reference:

https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store

NEW QUESTION: 60

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1.

You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

* Minimize how long it takes to complete the procedure.
* Minimize data loss.

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

Answer Area

Answer:

Answer Area

NEW QUESTION: 61

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. Always Encrypted with secure enclaves
D. Transparent Data Encryption (TDE)
Answer: (SHOW ANSWER)

NEW QUESTION: 62
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: (SHOW ANSWER)
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.
Incorrect Answers:

A: High - The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds

NEW QUESTION: 63
What should you implement to meet the disaster recovery requirements for the PaaS solution?
A. Availability Zones
B. failover groups
C. Always On availability groups
D. geo-replication
Answer: B (LEAVE A REPLY)
Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.
The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

NEW QUESTION: 64
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies that several tables are missing indexes.
You need to ensure that indexes are created for the tables.
What should you do?
A. RuntheDBCC SQLPERF command.
B. Run the dbcc dbreindex command.
C. Modify the automatic tuning settings for db1.
D. Modify the Query Store settings for db1.
Answer: (SHOW ANSWER)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview
NEW QUESTION: 65
Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.
You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the

database.
What should you use?
A. Azure Migrate
B. Data Migration Assistant (DMA)
C. Azure Data BOX
D. Azure Database Migration Service
Answer: (SHOW ANSWER)
https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Migrate

NEW QUESTION: 66
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 (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15

NEW QUESTION: 67
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.
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.

Answer:

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables

NEW QUESTION: 68

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.

You need to ensure that SOL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs.

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.

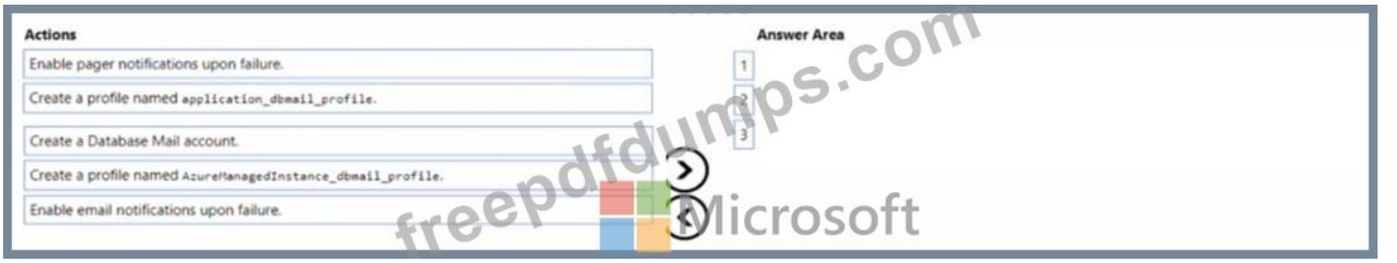

Answer:

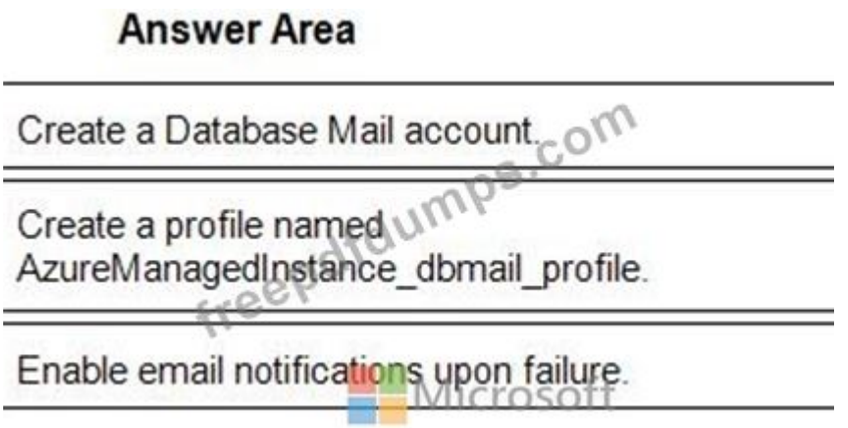

1 - Create a Database Mail account.

2 - Create a profile named AzureManagedInstance_dbmail_profile.

3 - Enable email notifications upon failure.

NEW QUESTION: 69

You have SQL Server 2019 or an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

Vou scale up the virtual machine to 8 vCPUs and 64 GB of memory.

You need to reduce tempdb contention without regatively affecting server performance.

What is the number of secondary data files that you should configure for tempdb?

A. 4

B. 2

C. 8

D. 64

Answer: C (LEAVE A REPLY)

NEW QUESTION: 70

You plan to migrate on-premises Microsoft SQL Server databases to Azure.

You need to identify which deployment and resiliency options meet the following requirements:

Support user-initiated backups.

Support multiple automatically replicated instances across Azure regions.

Minimize administrative effort to implement and maintain business continuity.

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

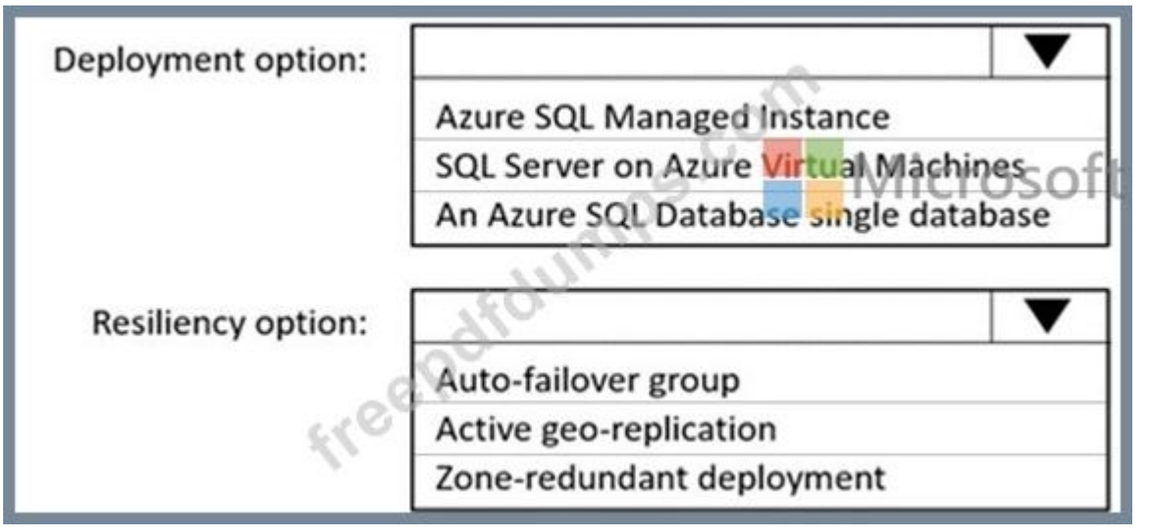

Answer:

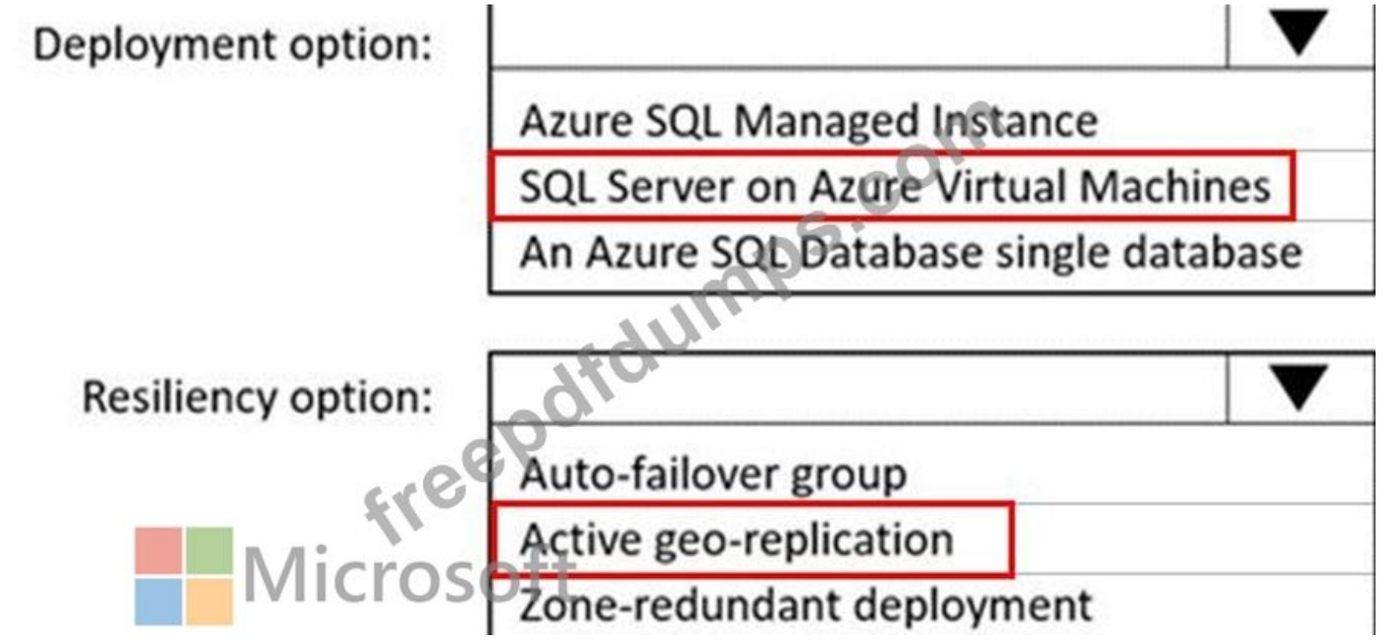

Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-is-overview
https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure

NEW QUESTION: 71
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 (LEAVE A REPLY)
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: 72
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. Microsoft SQL Server failover clusters
C. auto-failover groups
D. active geo-replication
Answer: D (LEAVE A REPLY)

NEW QUESTION: 73
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 (LEAVE A REPLY)
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: 74
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 (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/private-endpoint-overview

NEW QUESTION: 75
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 the Remove-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabase PowerShell cmdlet for Database1 on Server2.
Does this meet the goal?
A. Yes
B. No
Answer: B (LEAVE A REPLY)
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: 76
You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

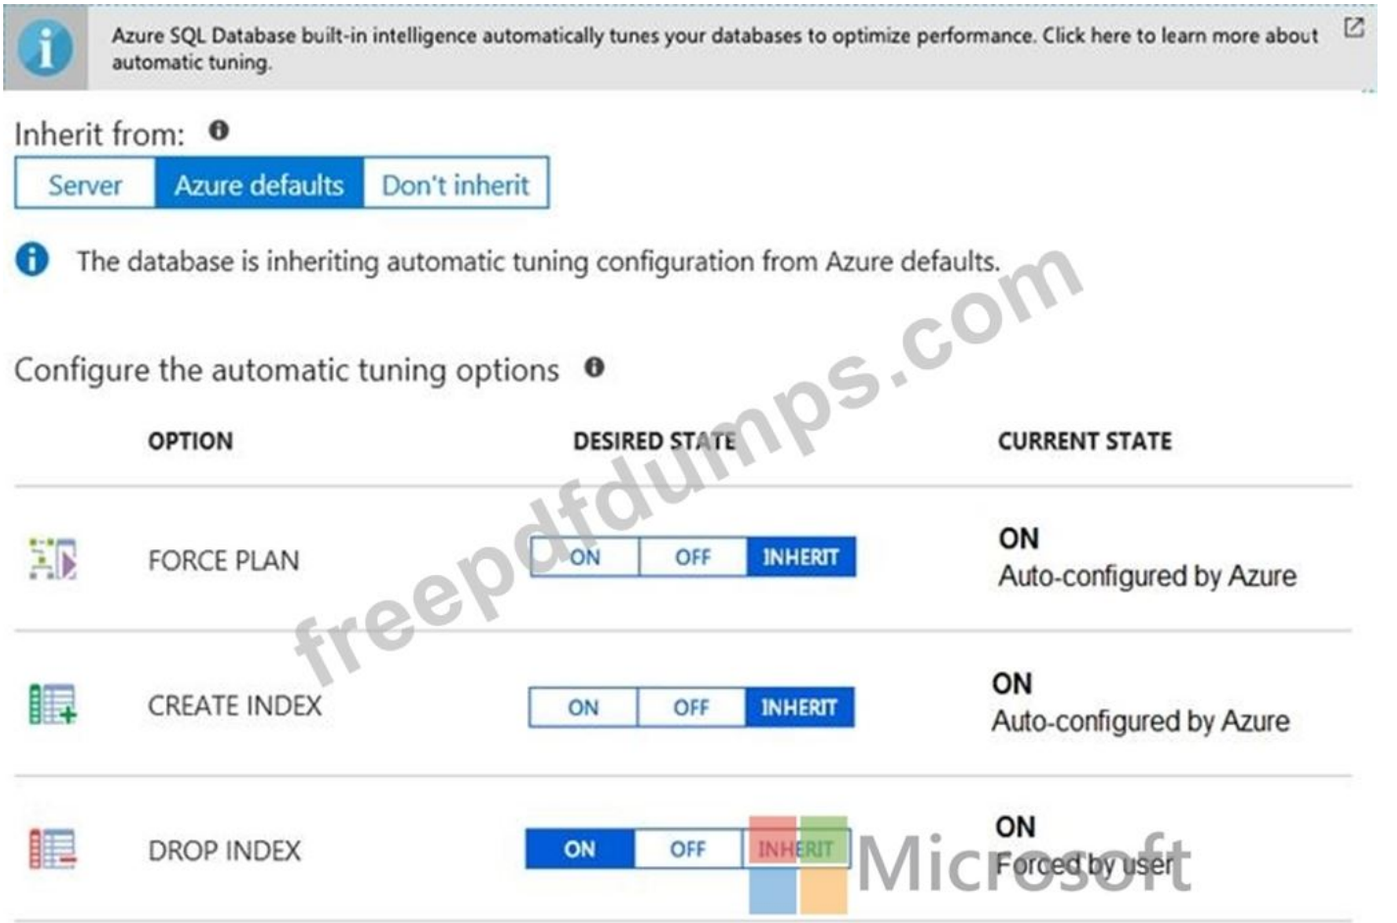

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Nonclustered indexes will be added to tables to improve performance.	○	○
Columns will be added to existing indexes automatically.	○	○
The query execution plan will revert to a previous plan if query performance degrades.	○	○

Answer:

NEW QUESTION: 77

You have an Azure subscription that contains the resources shown in the following table.

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

App1 experiences transient connection errors and timeouts when it attempts to access db1 after extended periods of inactivity.

You need to modify db1 to resolve the issues experienced by Appl as soon as possible, without considering immediate costs What you do?

A. Increase the number of vCores allocated to db1.

B. Enable automatic tuning for db1.

C. Disable auto-pause delay for db1.

D. Decrease the auto-pause delay for db1.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 78

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.	○	○

Answer:

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.	○	○

Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

NEW QUESTION: 79

You have an Azure SQL database named that contains a table named Table1.

You run a query to bad data into Table1.

The performance Of Table1 during the load operation are shown in exhibit.

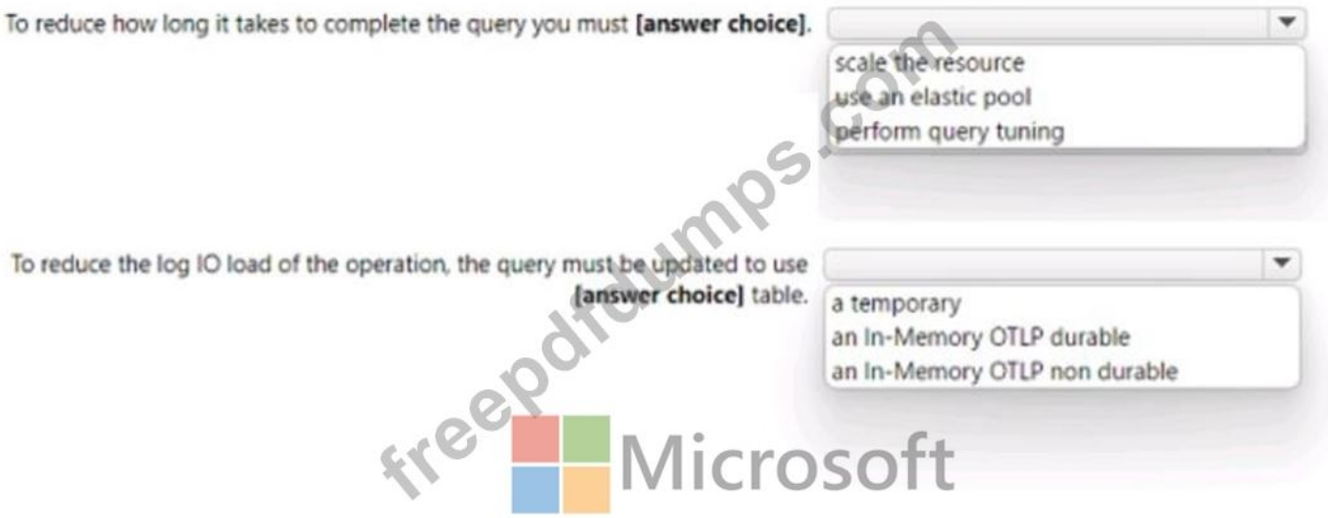

Answer:

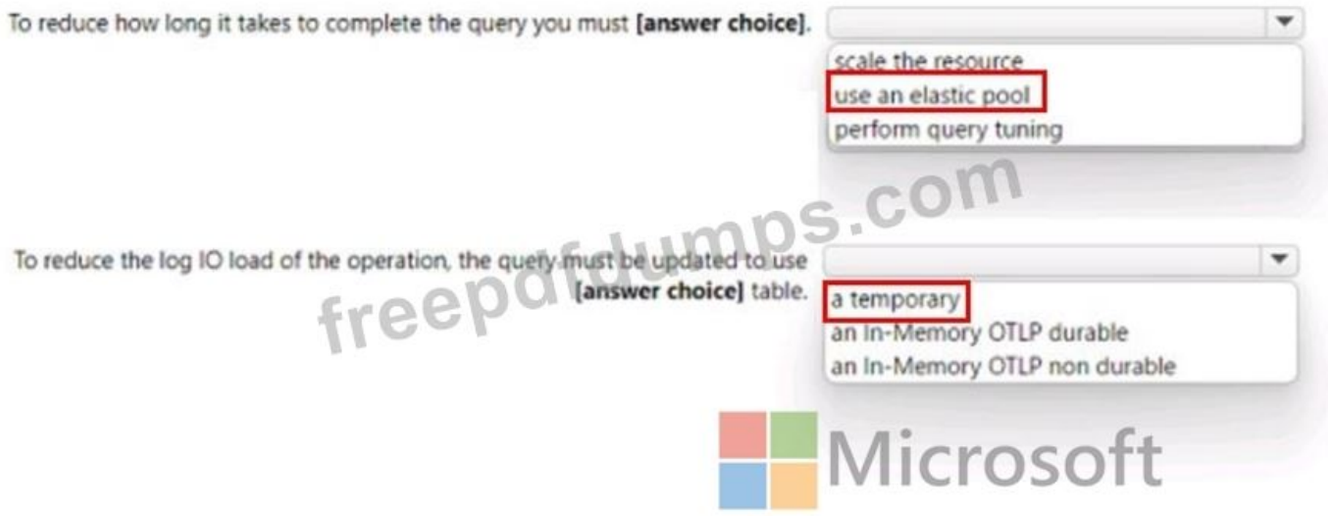

NEW QUESTION: 80

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 (LEAVE A REPLY)

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: 81
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.
You scale up the virtual machine to 16 vCPUSs and 64 GB of memory.
You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?
A. 2
B. 4
C. 8
D. 64
Answer: D (LEAVE A REPLY)
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

NEW QUESTION: 82
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 (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry

NEW QUESTION: 83

You have an Azure SQL database named DB1.
A user named User 1 has an Azure AD account.
You need to provide User1 with the ability to add and remove columns from the tables m DBV
The solution must use the principle of least privilege.
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 contained database user.
B. Assign the database user the db.ddladmm role.
C. Assign the database user the db.owner role.
D. Create a login and an associated database user.
Answer: (SHOW ANSWER)

NEW QUESTION: 84
You are monitoring an Azure Stream Analytics job.
You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero.
You need to ensure that the job can handle all the events.
What should you do?
A. Remove any named consumer groups from the connection and use $default.
B. Change the compatibility level of the Stream Analytics job.
C. Create an additional output stream for the existing input stream.
D. Increase the number of streaming units (SUs).
Answer: (SHOW ANSWER)
Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring

NEW QUESTION: 85
You have SQL Server on an Azure virtual machine.
You need to add a 4-TB volume that meets the following requirements:
Maximizes IOPs
Uses premium solid state drives (SSDs)
What should you do?
A. Attach two mirrored 4-TB SSDs.
B. Attach a stripe set that contains four 1-TB SSDs.
C. Attach a RAID-5 array that contains five 1-TB SSDs.
D. Attach a single 4-TB SSD.
Answer: (SHOW ANSWER)
Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=windows2016

NEW QUESTION: 86
You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently.
You enable columnstore archival compression.
What are two possible results of the configuration? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.
A. Queries that use the index will consume more CPU resources.
B. Queries that use the index will consume more disk I/O.
C. The index will consume more disk space.
D. The index will consume more memory.
E. Queries that use the index will retrieve fewer data pages.
Answer: (SHOW ANSWER)
For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.
Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

NEW QUESTION: 87
You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1.
You need to configure the autogrow and autoshrink settings for DB1.
Which statements should you use? To answer, select the appropriate options in the answer are a.
NOTE: Each correct selection is worth one point.

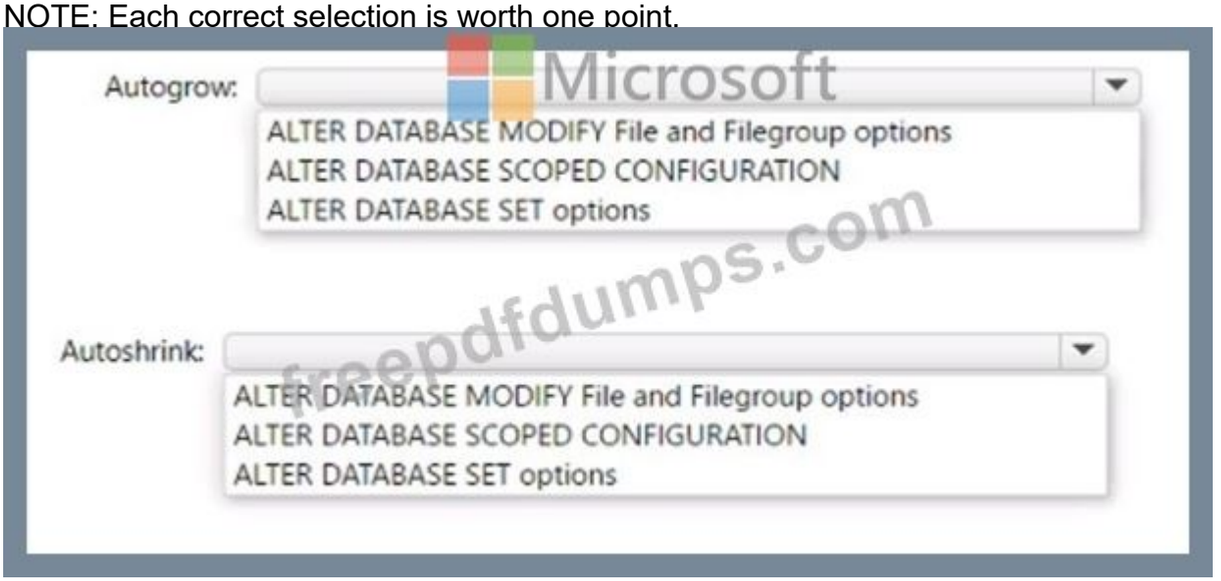

Answer:

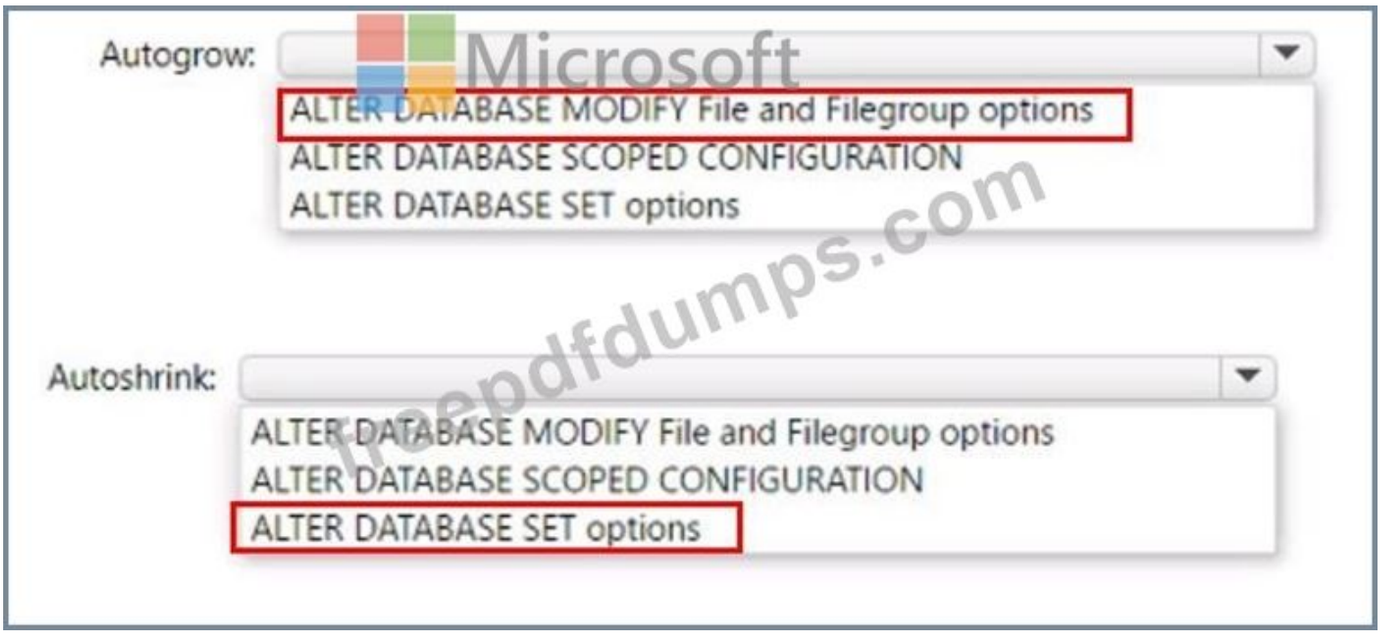

NEW QUESTION: 88
You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following 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 (LEAVE A REPLY)
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: 89
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), run sp_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 (LEAVE A REPLY)
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-database-mail

NEW QUESTION: 90
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: (SHOW ANSWER)
https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/

NEW QUESTION: 91
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 (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

NEW QUESTION: 92
You have an Azure subscription.
Vou need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:
* Minimize how long it takes to fail over.
* Maintain existing connections to the primary replica during a failover.
What should you do?
A. Connect each virtual machine to a different subnet on a single virtual network.
B. Connect each virtual machine to a single subnet on a single virtual network.
C. Connect each virtual machine to a different subnet on a virtual network. Deploy a basic Azure load balancer.
D. Connect each virtual machine to a single subnet on a virtual network. Deploy a standard Azure load balancer.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 93
You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.
You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.
You need to recommend a database service for the deployment. The solution must minimize administrative effort.
What should you include in the recommendation?
A. Azure SQL Managed instance
B. Azure SQL Database single database
C. Azure Database for PostgreSQL
D. SQL Server on Azure virtual machines
Answer: B (LEAVE A REPLY)
Azure SQL Database single database supports Data Sync.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison

NEW QUESTION: 94
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: (SHOW ANSWER)
Reference:
https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/ This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is

independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
Topic 3, ADatum Corporation
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.
Existing Environment
ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.
SALESDB collects data from the stores and the website.
DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.
REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.
Requirements
Planned Changes
ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:
Migrate SALESDB and REPORTINGDB to an Azure SQL database.
Migrate DOCDB to Azure Cosmos DB.
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.
As they arrive, all the sales documents in JSON format must be transformed into one consistent format.
Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.
Technical Requirements
The new Azure data infrastructure must meet the following technical requirements:
Data in SALESDB must encrypted by using Transparent Data Encryption (TDE). The encryption

must use your own key.
SALESDB must be restorable to any given minute within the past three weeks.
Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.
Missing indexes must be created automatically for REPORTINGDB.
Disk IO, CPU, and memory usage must be monitored for SALESDB.

NEW QUESTION: 95
You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.
During peak usage, the database will require the following:
24 cores
500 GB of storage
124 GB of memory
More than 50,000 IOPS
During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard.
Which service tier should you use during peak usage?
A. Business Critical
B. Premium
C. Hyperscale
Answer: (SHOW ANSWER)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcore-single-databases#business-critical---provisioned-compute---gen4

NEW QUESTION: 96
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.
You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

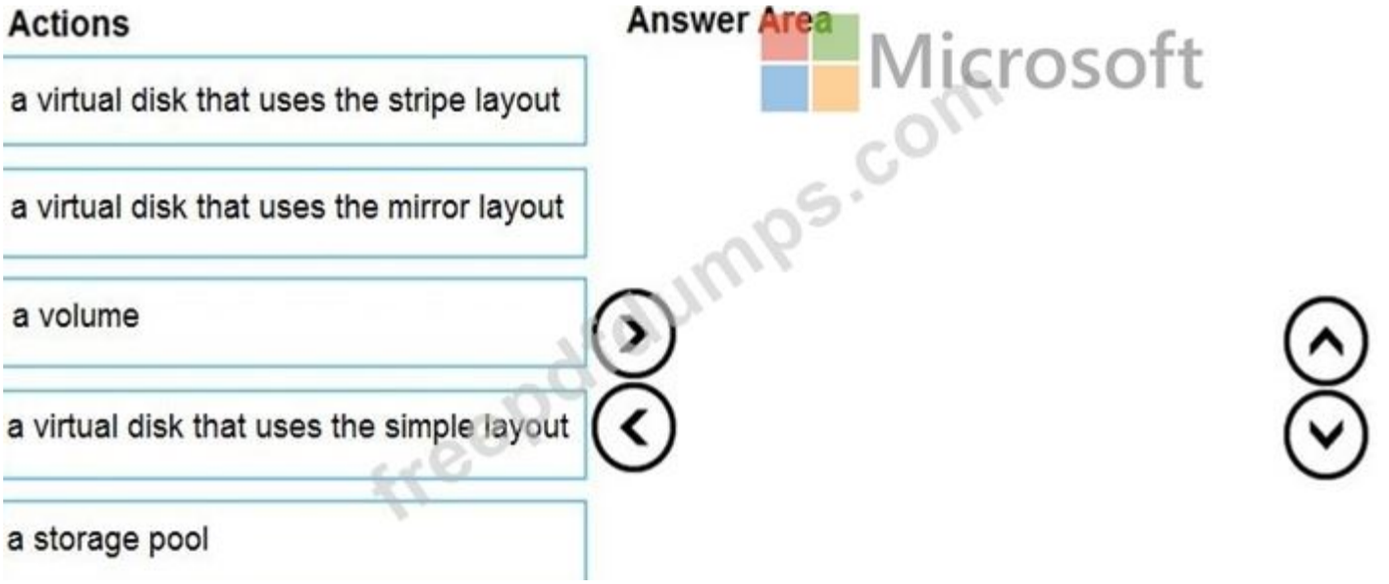

Answer:

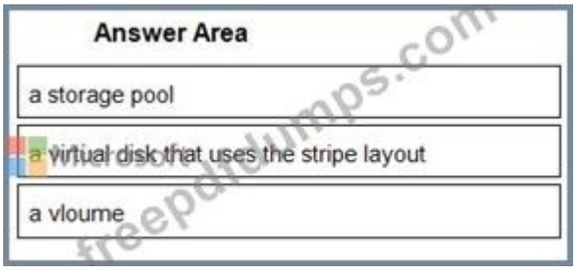

1 - a storage pool

2 - a virtual disk that uses the stripe layout

3 - a vloume

Reference:

https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/

NEW QUESTION: 97

You have an Azure Data Factory that contains 10 pipelines.

You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.

What should you add to each pipeline?

A. an annotation

B. a resource tag

C. a run group ID

D. a user property

E. a correlation ID

Answer: A (LEAVE A REPLY)

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.

Reference:

https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/

NEW QUESTION: 98

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

You need to recommend changes to the data types to reduce storage and improve performance.

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

NOTE: Each correct selection is worth one point.

A. Change PhoneNumber to the bigint data type.

B. Change PhoneNumber to the float data type.

C. Change LastHireDate to the date data type.

D. Change Salary to the money data type.

E. Change LastHireDate to the datetime2(7) data type.

Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 99

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

Send the output to an Azure Synapse.

Identify spikes and dips in time series data.

Minimize development and configuration effort.

Which should you include in the solution?

A. Azure SQL Database

B. Azure Databricks

C. Azure Stream Analytics

Answer: C (LEAVE A REPLY)

Anomalies can be identified by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics Reference:

https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/

https://docs.microsoft.com/en-us/azure/stream-analytics/azure-synapse-analytics-output

NEW QUESTION: 100

You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.
You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication.
You need to ensure that App1 can access DB2. The solution must meet the following requirements:
App1 must be able to view only DB2.
Administrative effort must be minimized.
What should you create?
A. a contained database user for App1 on DB2
B. a login for App1 on Server1
C. a contained database user from an external provider for App1 on DB2
D. a contained database user from a Windows login for App1 on DB2
Answer: D (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-database-portable?view=sql-server-ver15

NEW QUESTION: 101
You are evaluating the business goals.
Which feature should you use to provide customers with the required level of access based on their service agreement?
A. dynamic data masking
B. Conditional Access in Azure
C. service principals
D. row-level security (RLS)
Answer: (SHOW ANSWER)
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15

NEW QUESTION: 102
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.

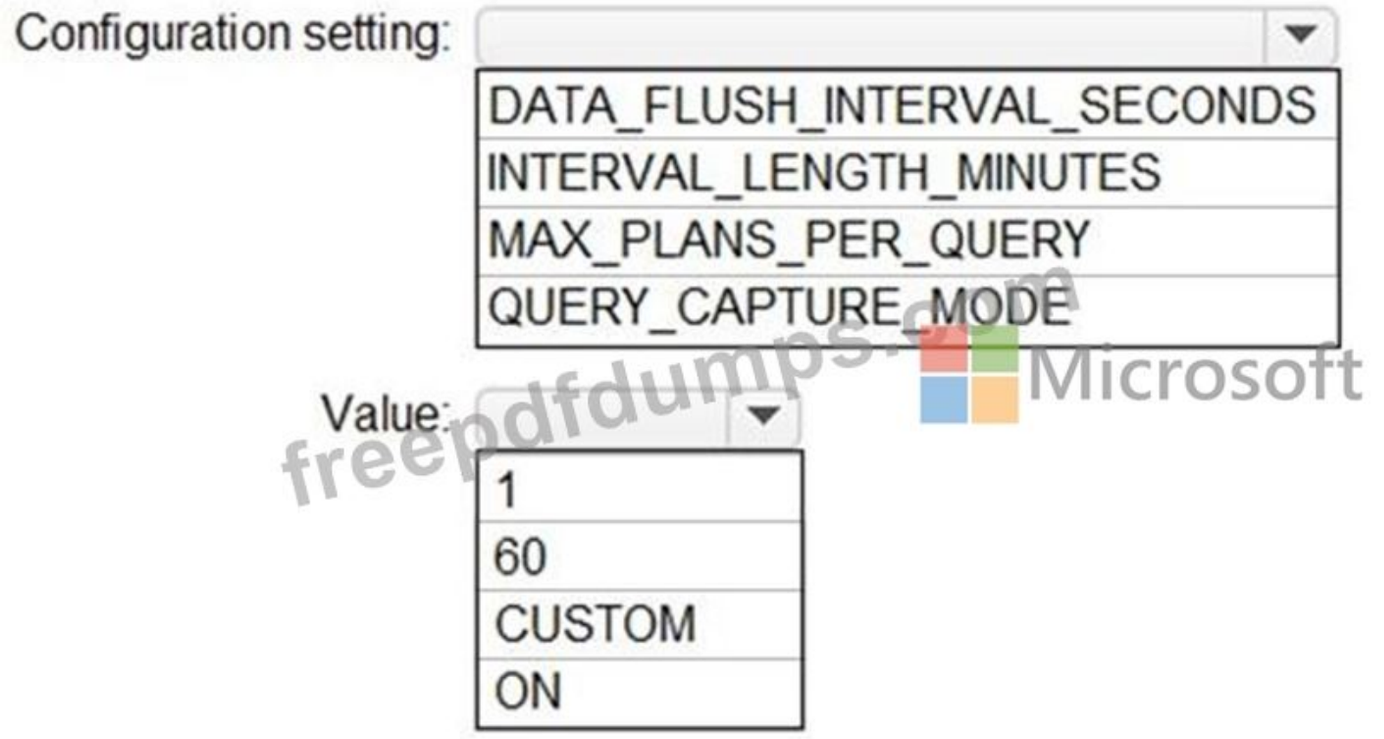

Answer:

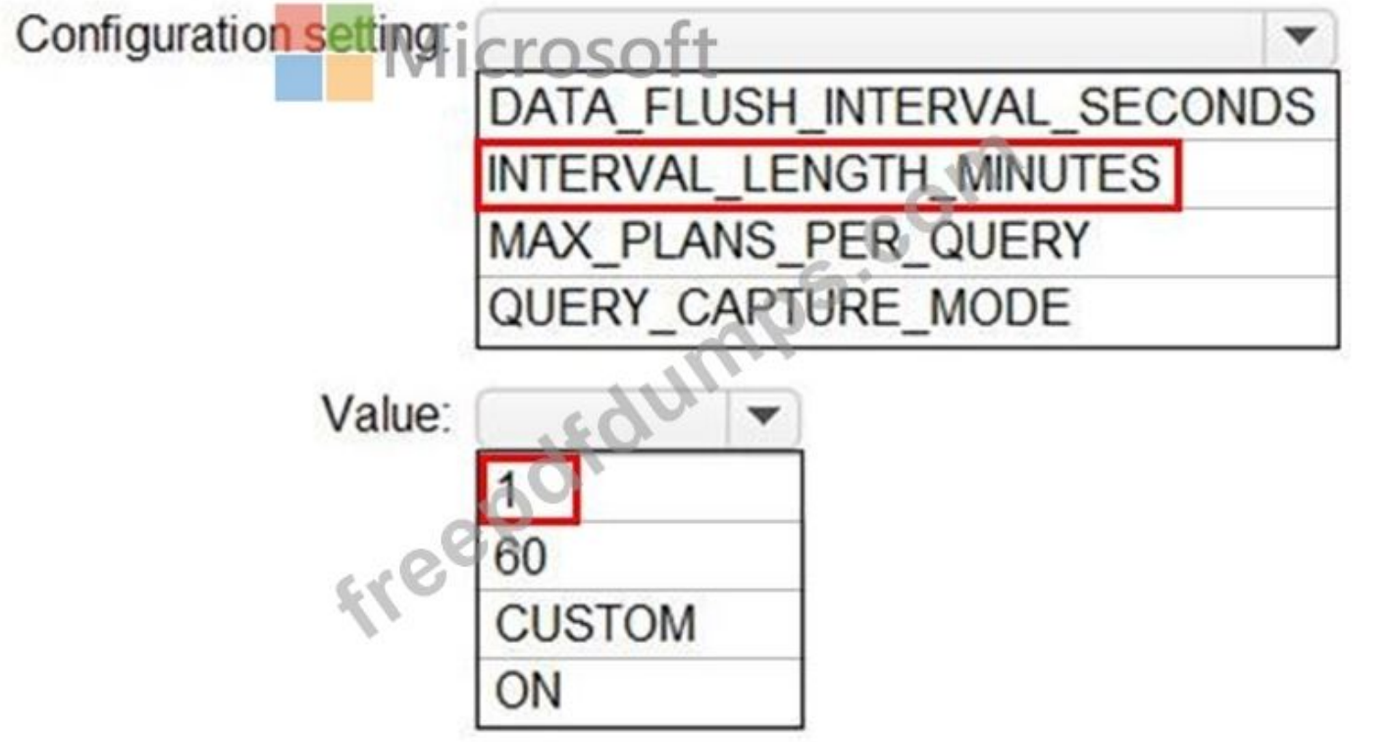

NEW QUESTION: 103

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.

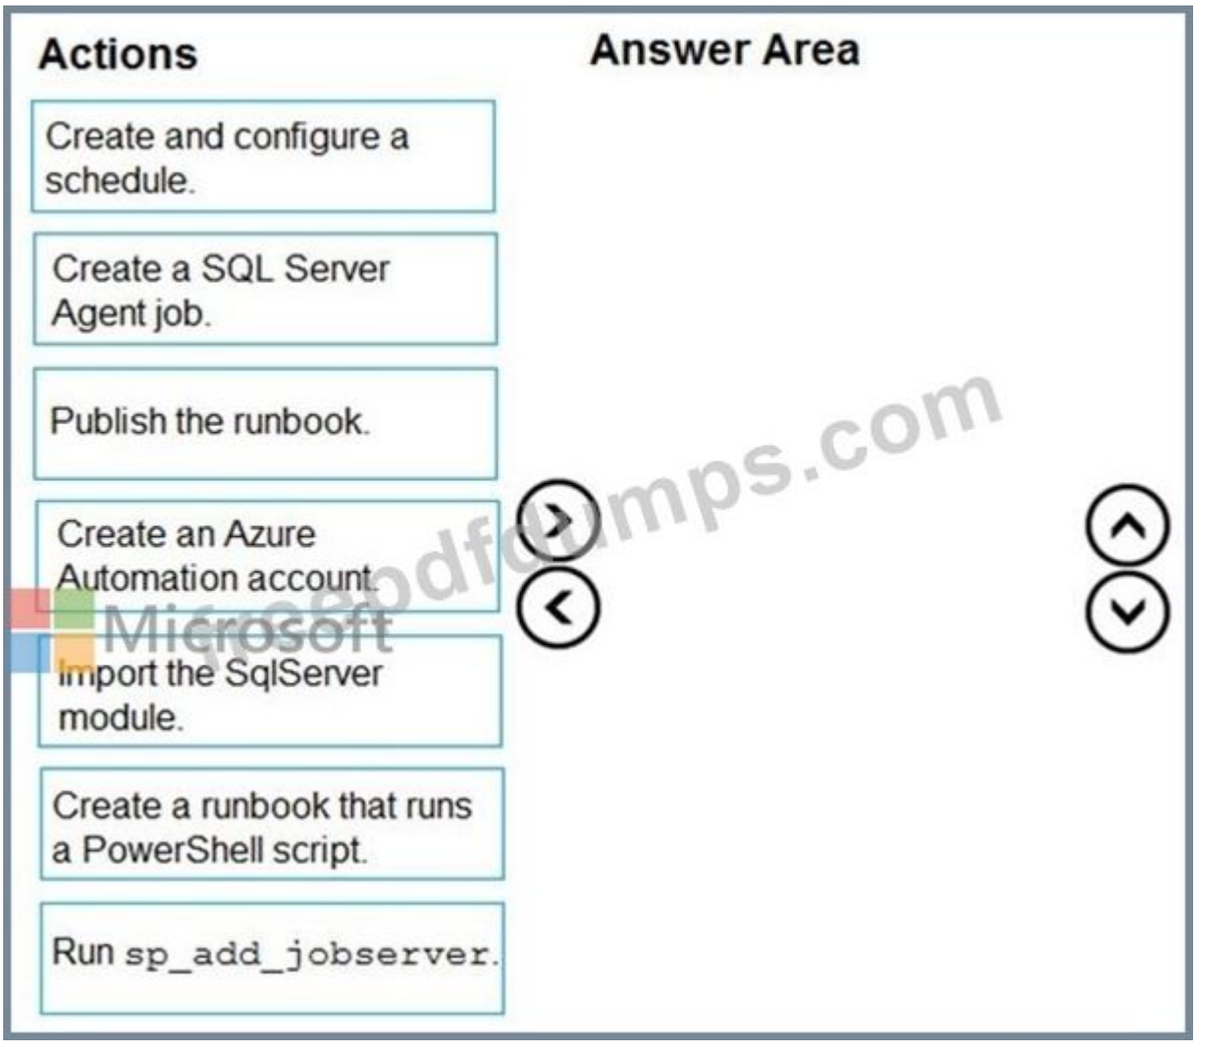

Answer:

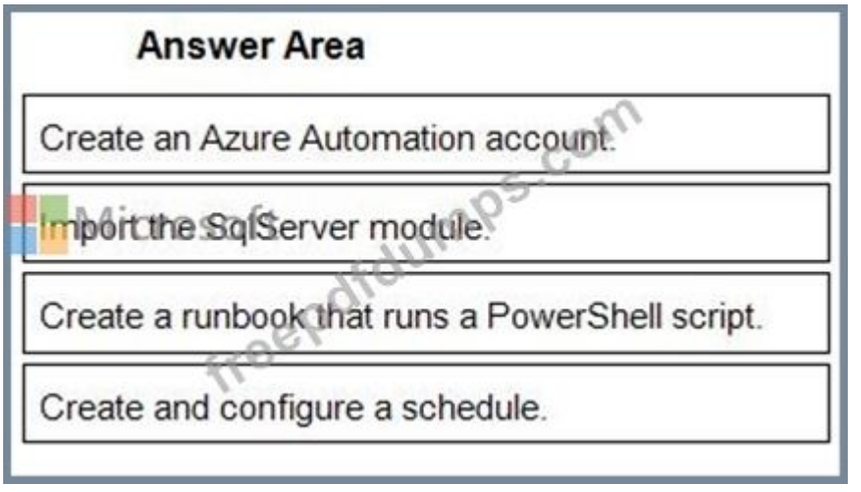

1 - Create an Azure Automation account.

2 - Import the SqlServer module.

3 - Create a runbook that runs a PowerShell script.

4 - Create and configure a schedule.

Reference:

https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statistics-maintenance-using/ba-p/368974

NEW QUESTION: 104
You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order.
You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders.
What should you include in the solution? To answer, select the appropriate options in the answer are a.
NOTE: Each correct selection is worth one point.

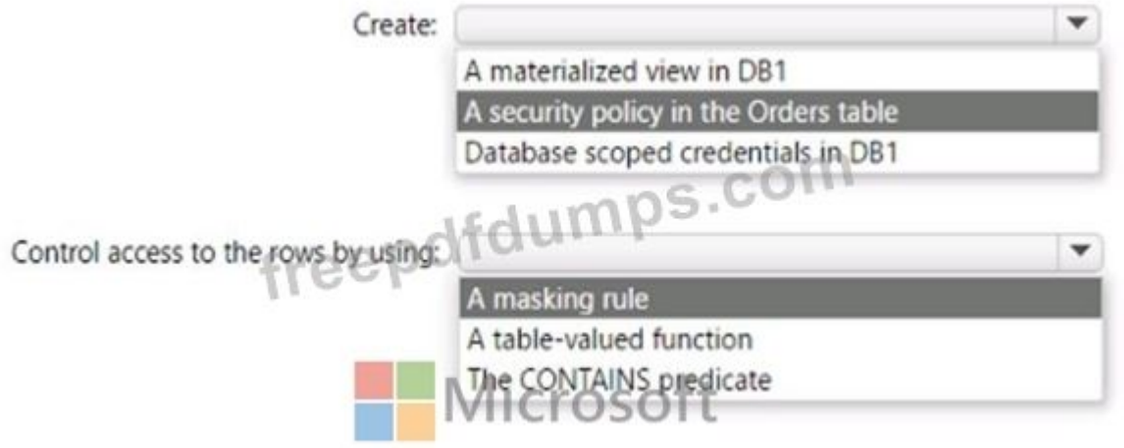

Answer:

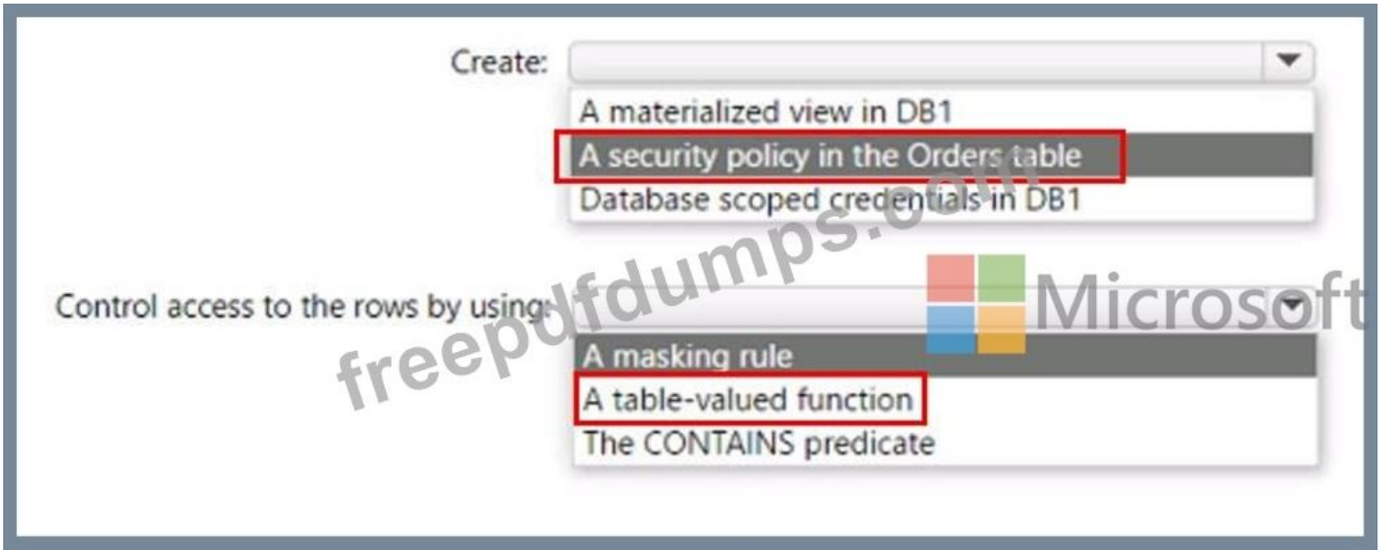

NEW QUESTION: 105
You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2 . SQLM2 contains a database named DB1 and a user named User1.
User1 drops DB1.
You need to perform a point-in-time restore of DB1 to SQLMI2.

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 106

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:

The first six numerals of each customer's phone number must be masked.

The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

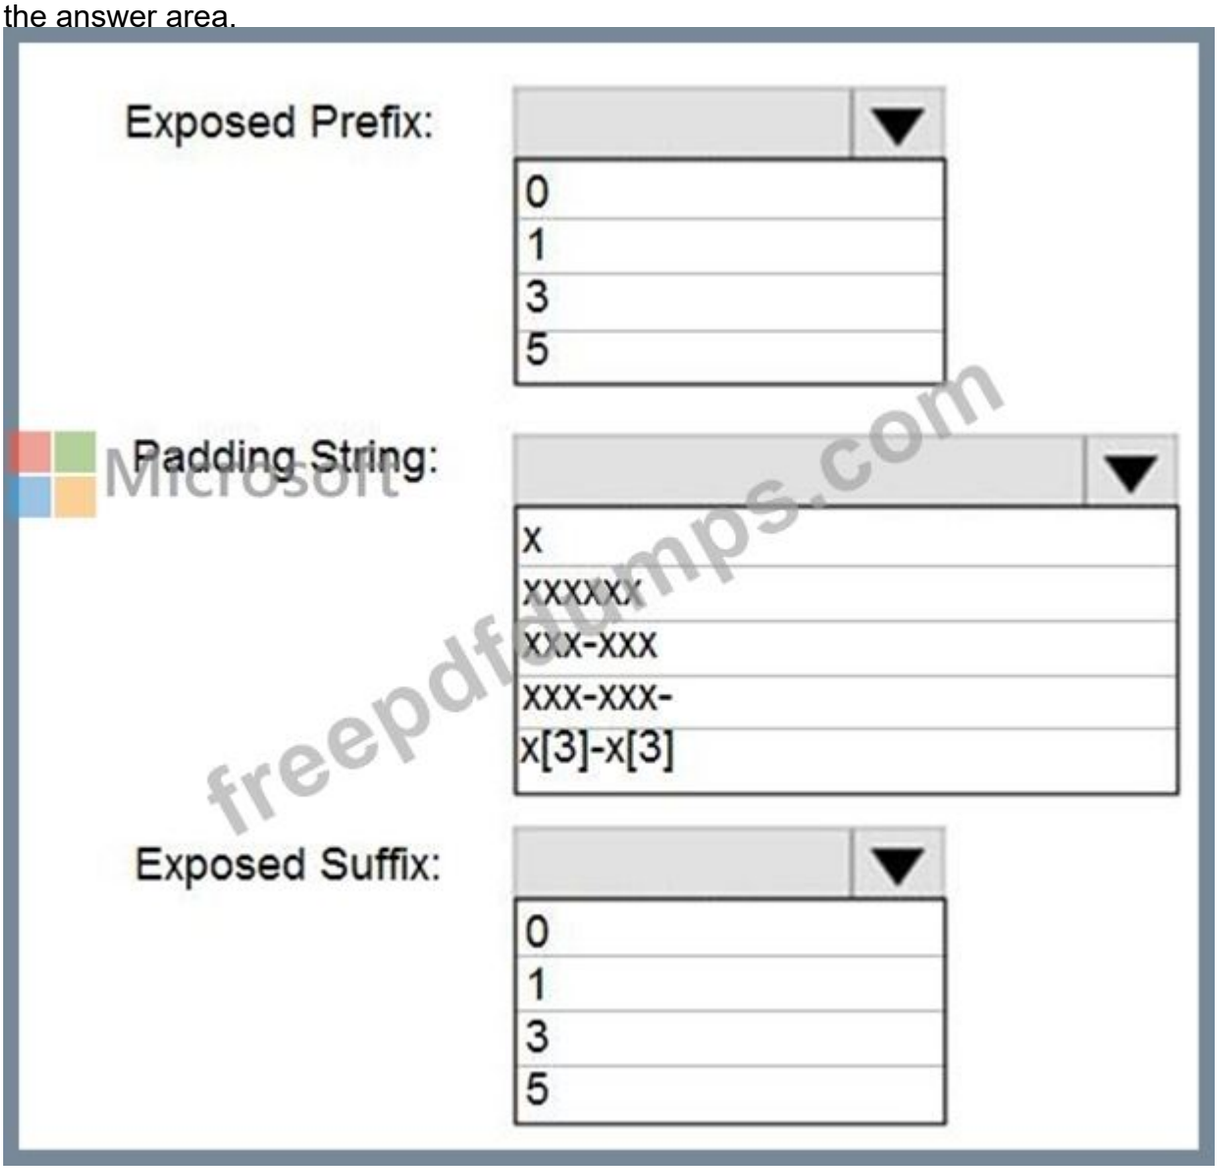

Answer:

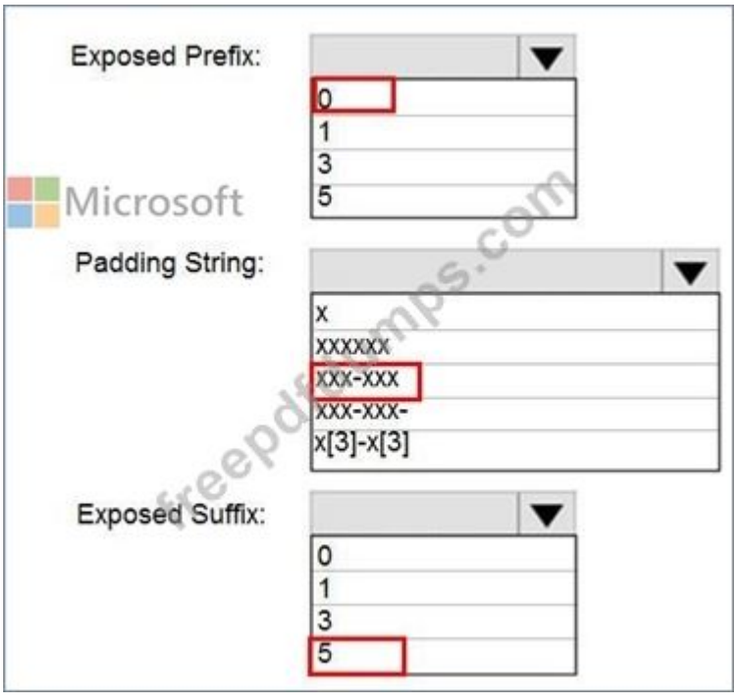

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

NEW QUESTION: 107
You have an Azure SQL database named DB 1 in the General Purpose service tier.
You need to monitor DB 1 by using SQL Insights.
What should you include in the solution? To answer, select the appropriate options in the answer are a.
NOTE: Each correct selection is worth one point.

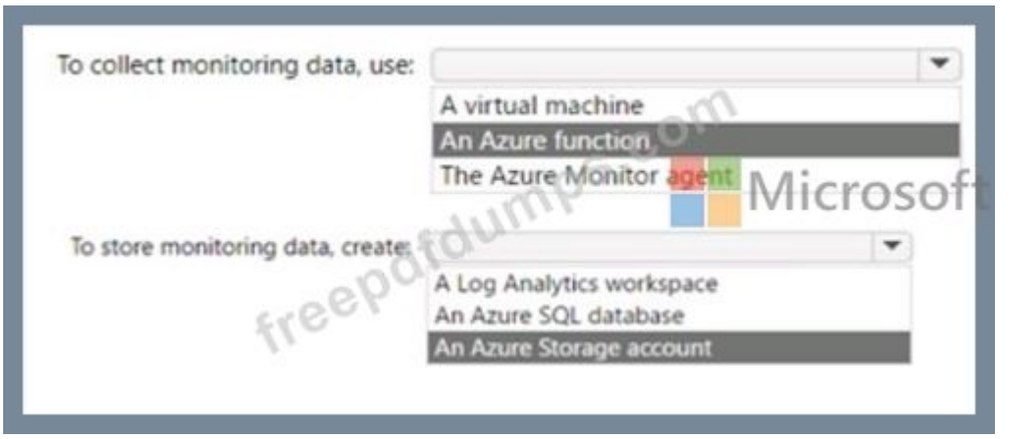

Answer:

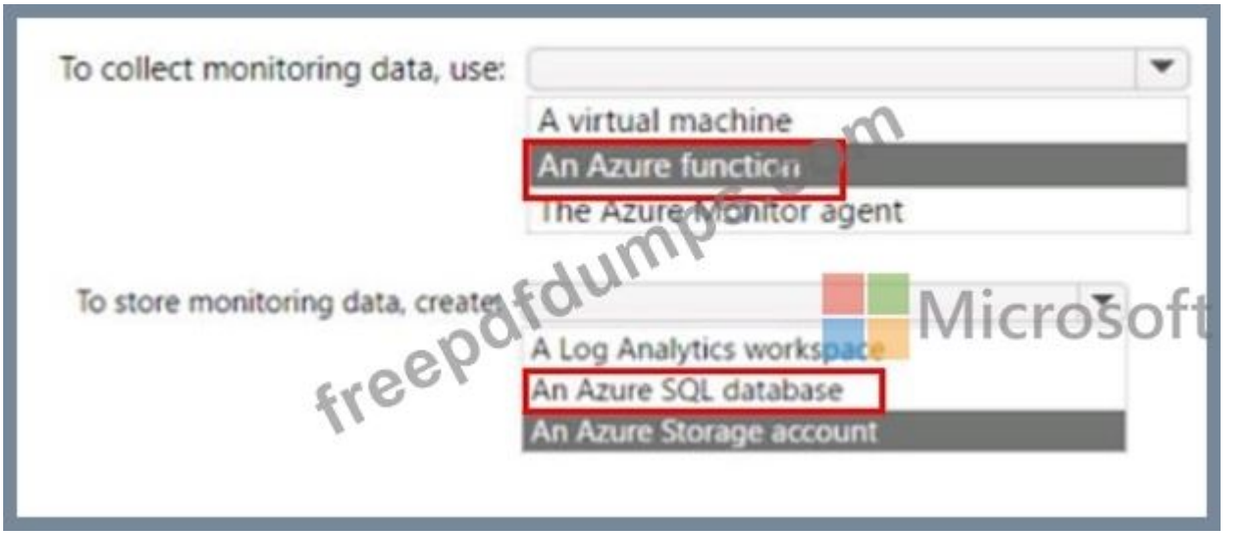

NEW QUESTION: 108

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.

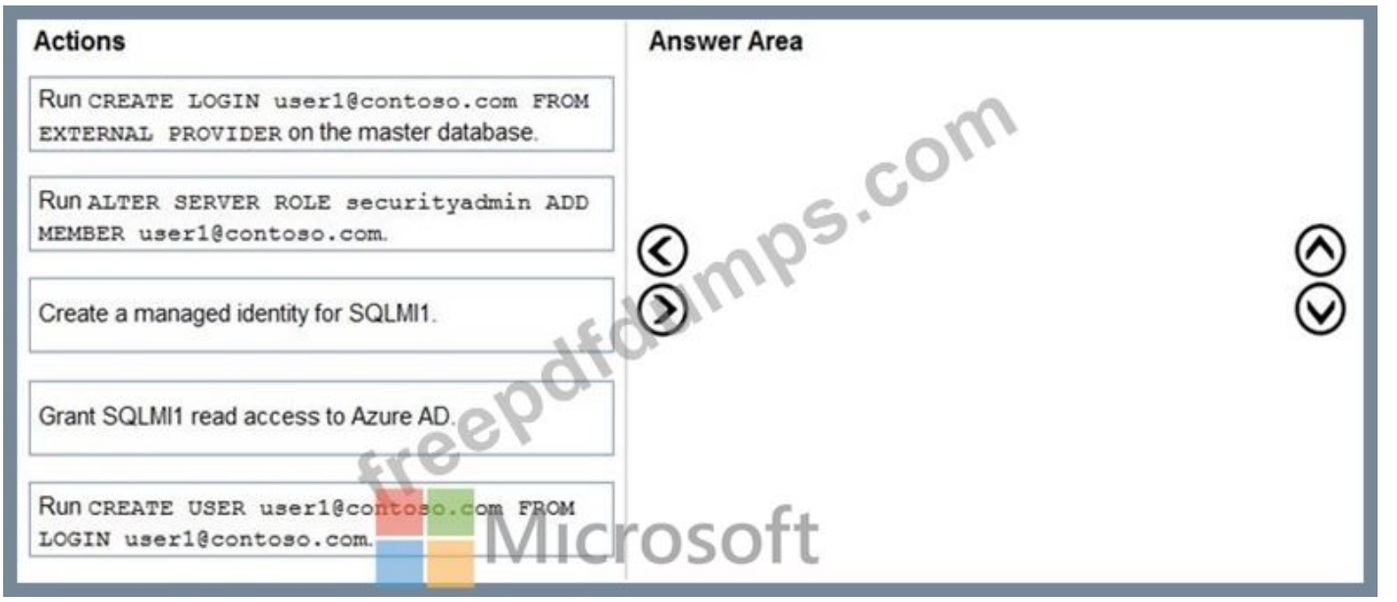

Answer:

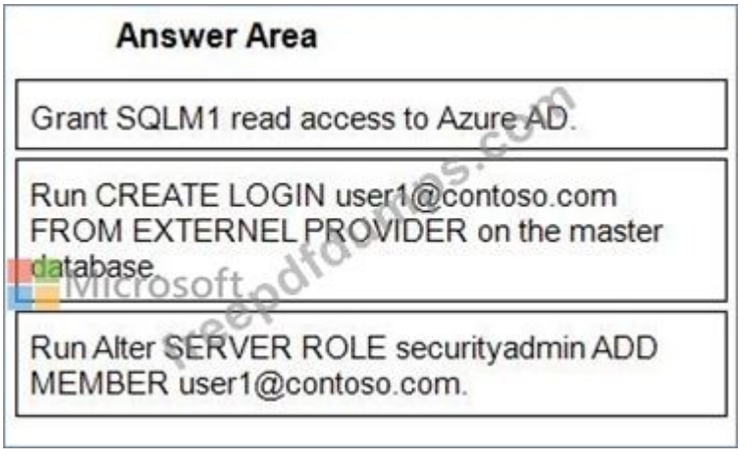

1 - Grant SQLM1 read access to Azure AD.

2 - Run CREATE LOGIN user1@contoso.com FROM EXTERNEL PROVIDER on the master database.

3 - Run Alter SERVER ROLE securityadmin ADD MEMBER user1@contoso.com.

Reference:

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

NEW QUESTION: 109

Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1.

You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance.

Which service tier should you use?

A. Premium

B. Business Critical

C. Hyperscale

D. General Purpose

Answer: B (LEAVE A REPLY)

NEW QUESTION: 110

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: (SHOW ANSWER)

Reference:

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

NEW QUESTION: 111

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.

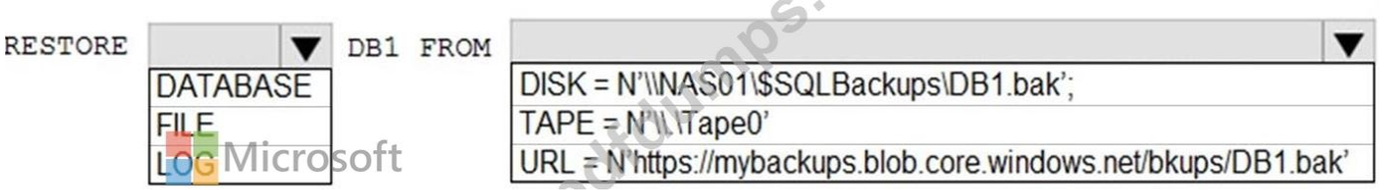

Answer:

NEW QUESTION: 112

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 (LEAVE A REPLY)
Reference:
https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance

NEW QUESTION: 113
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 SQL! to use mixed mode authentication. Which procedure should you run?
A. xp_instance_regwrite
B. sp_cnarge_users_login
C. sp_addremotelogin
D. xp_grant_login
Answer: A (LEAVE A REPLY)

NEW QUESTION: 114
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 dat a. 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 an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
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: 115
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 (LEAVE A REPLY)
DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database.
Incorrect Answers:
D: To shrink one data or log file at a time for a specific database, execute the DBCC SHRINKFILE command.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql

NEW QUESTION: 116
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 (LEAVE A REPLY)
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: 117
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 the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.
Does this meet the goal?
A. Yes

B. No

Answer: B (LEAVE A REPLY)

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