

Microsoft.DP-300.v2026-09-14.q168

Exam Code:	DP-300
Exam Name:	Administering Relational Databases on Microsoft Azure
Certification Provider:	Microsoft
Free Question Number:	168
Version:	v2026-09-14
# of views:	102
# of Questions views:	1723
https://www.freecram.net/torrent/Microsoft.DP-300.v2026-09-14.q168.html	

NEW QUESTION: 1

You have an Azure subscription that contains an Azure SQL database named DB1.
You need to host elastic jobs by using DB1. DB1 will also be configured as a job target. The solution must support the use of location-based Conditional Access policies.
What should the elastic jobs use to access DB1?

- A. a system-assigned managed identity
- B. Azure SQL sign-in credentials
- C. database-scoped credentials
- D. a user-assigned managed identity

Answer: (SHOW ANSWER)

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.
Why C is Correct: database-scoped credentials can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. The scenario asks for: You need to host elastic jobs by using DB1. That makes database-scoped credentials the option that satisfies the required Azure SQL layer and operational outcome.
Why A is Wrong: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. It would leave too much manual work or would operate at the wrong scope for the requested automation.
Why B is Wrong: Azure SQL sign-in credentials can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.
Why D is Wrong: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. It would leave too much manual work or would operate at the wrong scope for the requested automation.

NEW QUESTION: 2

You plan to develop a dataset named Purchases by using Azure Databricks. Purchases will contain the following columns:
ProductID
ItemPrice
LineTotal
Quantity
StoreID
Minute

Month
Hour
Year
Day

You need to store the data to support hourly incremental load pipelines that will vary for each StoreID. The solution must minimize storage costs.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

df.write

.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour", "StoreID")

.mode ("append")

.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

Answer:

df.write

.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour", "StoreID")

.mode ("append")

.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

Explanation:

df.write

.bucketBy	("*")
.partitionBy	("StoreID", "Hour")
.range	("StoreID", "Year", "Month", "Day", "Hour")
.sortBy	("Year", "Month", "Day", "Hour", "StoreID")

.mode ("append")

.csv("/Purchases")
.json("/Purchases")
.parquet("/Purchases")
.saveAsTable("/Purchases")

Box 1: .partitionBy

Example:

```
df.write.partitionBy( " y " , " m " , " d " )
mode(SaveMode.Append)
parquet( " /data/hive/warehouse/db_name.db/ " + tableName)
```

Box 2: (" Year " , " Month " , " Day " , " Hour " , " StoreID ")

Box 3: .parquet(" /Purchases ")

Reference:

<https://intellipaat.com/community/11744/how-to-partition-and-write-dataframe-in-spark-without-deleting-partitions-with-no-new-data> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: Purchases will contain the following columns: ProductID ItemPrice LineTotal Quantity StoreID Minute Month Hour Year Day You need to store the data to support hourly incremental load pipelines that will vary for each StoreID.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 3

You have two instances of SQL Server on Azure virtual Machines named VM1 and VM2. VM1 hosts a database named db1.

You plan to create a database availability group (DAG) for db1. The solution must use certificate authentication between VM1 and VM2.

You need to configure authentication for the outbound connections of VM1.

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

Actions

From the master database, run CREATE ENDPOINT.

From db1, run CREATE CERTIFICATE.

From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.

From the master database, run CREATE CERTIFICATE.

From db1, run CREATE ENDPOINT.

Answer Area

Answer:

Actions

From the master database, run CREATE ENDPOINT.

From db1, run CREATE CERTIFICATE.

From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.

From the master database, run CREATE CERTIFICATE.

From db1, run CREATE ENDPOINT.

Answer Area

From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.

From the master database, run CREATE CERTIFICATE.

From db1, run CREATE ENDPOINT.

Explanation:

Actions

From the master database, run CREATE ENDPOINT.

From db1, run CREATE CERTIFICATE.

>

<

Answer Area

1From the master database, run CREATE MASTER KEY ENCRYPTION BY PASSWORD.

2From the master database, run CREATE CERTIFICATE.

3From db1, run CREATE ENDPOINT.

⬆

⬇

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: The solution must use certificate authentication between VM1 and VM2.

Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

NEW QUESTION: 4

You have an Azure SQL database named DB1.

You need to encrypt DB1. The solution must meet the following requirements;

- * Encrypt data in motion.
- * Support comparison operators.
- * Provide randomized encryption.

What should you include in the solution?

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why D is Correct: Always Encrypted protects sensitive columns from exposure to the database engine by encrypting and decrypting data at the client side with keys protected outside the database. In this scenario, the important constraint is: You need to encrypt DB1. Always Encrypted with secure enclaves satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Always Encrypted protects sensitive columns from exposure to the database engine by encrypting and decrypting data at the client side with keys protected outside the database. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: column-level encryption is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to encrypt DB1.

Why C is Wrong: Transparent Data Encryption protects database files, backups, and logs at rest; it does not encrypt data in use or enforce user-level data visibility. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 5

You have an Azure subscription. The subscription contains an instance of SQL Server on Azure Virtual Machines named SQL1 and an Azure Automation account named account1. You need to configure account1 to restart the SQL Server Agent service if the service stops. Which setting should you configure?

- A. Start/Stop VM
- B. Update management
- C. State configuration (DSC)
- D. Change tracking

Answer: (SHOW ANSWER)

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why C is Correct: State configuration (DSC) can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. This fits the case because the requested outcome is: You need to configure account1 to restart the SQL Server Agent service if the service stops. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Start/Stop VM can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why B is Wrong: Update management can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why D is Wrong: Change tracking can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

NEW QUESTION: 6

You have an Azure subscription that contains an Azure SQL managed instance named SQLMi1 and a SQL Agent job named Backupdb. Backupdb performs a daily backup of the databases hosted on SQLMi1. You need to be notified by email if the 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

Create a SQL Server Agent alert.

Create an operator.

Create an extended event.

Enable Database Mail.

Add a failure notification to the job.

Answer Area

⬅

➡

⬆

⬇

Microsoft

Answer:

- Actions**
- Create a SQL Server Agent alert.
 - Create an operator.
 - Create an extended event.
 - Enable Database Mail.
 - Add a failure notification to the job.

- Answer Area**
- Enable Database Mail.
 - Create an operator.
 - Add a failure notification to the job.

Explanation:

Microsoft

Enable Database Mail.

Create an operator.

Add a failure notification to the job.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to be notified by email if the job fails.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 7

You have an Azure SQL database named DB1. DB1 has a table named Table1 that contains the following columns.

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

You plan to enable Always Encrypted for Table1.

Which two columns support encryption? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point

- A. Column1
- B. Column2

- C. Column3
- D. Column4
- E. Column5

Answer: (SHOW ANSWER)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why A and D is Correct: Column1 is correct because it is the feature whose normal purpose matches the stated requirement. Column1 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: Which two columns support encryption? Column4 is correct because it is the feature whose normal purpose matches the stated requirement. Column4 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: Which two columns support encryption?

Why B is Wrong: Column2 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: Which two columns support encryption?

Why C is Wrong: Column3 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why E is Wrong: Column5 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: Which two columns support encryption?

NEW QUESTION: 8

You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

Microsoft Azure

Search resources, services, and docs (G+)

liveid238@outlook.com
DEFAULT DIRECTORY

Home > SQL databases > Create SQL Database >

Configure

Feedback

Compute Hardware

Click "Change configuration" to see details for all hardware generations available including memory optimized and compute optimized options

Hardware Configuration

Gen5

up to 40 vCores, up to 120 GB memory

Change configuration

Max vCores

6 vCores

Min vCores

0.75 vCores

2.25 GB MIN MEMORY 18 GB MAX MEMORY

Auto-pause delay

The database automatically pauses if it is inactive for the time period specified here, and automatically resumes when database activity recurs. Alternatively, auto-pausing can be disabled.

☐ Enable auto-pause

DaysHoursMinutes

040

Data max size *

800 GB

240 GB LOG SPACE ALLOCATED

Apply

Cost summary

Gen5 - General Purpose (GP_5, Gen5_0)

Cost per GB (in USD)

0.12

Max storage selected (in GB)

x 1040

ESTIMATED STORAGE COST / MONTH

119.60 USD

COMPUTE COST / VCORE / SECOND ¹

0.000145 USD

NOTES

¹ Serverless databases are billed in vCores based on a combination of CPU and memory utilization. [Learn more about serverless billing](#)

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

NOTE: Each correct selection is worth one point.

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

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

 Microsoft

no extra time
up to 10 minutes
up to one minute

intermittent and unpredictable
regular and high
steady and low

Answer:

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

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

 Microsoft

no extra time
up to 10 minutes
up to one minute

intermittent and unpredictable
regular and high
steady and low

Explanation:

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

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

 Microsoft

no extra time
up to 10 minutes
up to one minute

intermittent and unpredictable
regular and high
steady and low

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You are provisioning an Azure SQL database in the Azure portal as shown in the following exhibit.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 9

You have an on-premises Microsoft SQL Server 2022 instance that hosts a database named DB1. You have an Azure subscription that contains an Azure SQL database named SQLDB1. You need to replicate DB1 to SQLDB1. Which type of replication should you use?

- A. snapshot
- B. peer-to-peer
- C. merge
- D. transactional

Answer: (SHOW ANSWER)

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: peer-to-peer is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. In this scenario, the important constraint is: You need to replicate DB1 to SQLDB1. peer-to-peer satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: snapshot is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: merge is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: transactional is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 10

You have SQL Server on an Azure virtual machine that contains a database named DB1.

The database reports a CHECKSUM error.

You need to recover the database.

How should you complete the statements? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

USE master;

ALTER DATABASE [DB1] SET

	▼
OFFLINE	
ONLINE	
SINGLE_USER	
TRUSTWORTHY	

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

	▼
MOINDEX	
PHYSICAL_ONLY	
REPAIR_ALLOW_DATA_LOSS	
REPAIR_FAST	

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

	▼
MULTI_USER;	
ONLINE;	
OPEN;	
TRUSTWORTHY;	

GO

Answer:

USE master;

ALTER DATABASE [DB1] SET

	▼
OFFLINE	
ONLINE	
SINGLE_USER	
TRUSTWORTHY	

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

	▼
MOINDEX	
PHYSICAL_ONLY	
REPAIR_ALLOW_DATA_LOSS	
REPAIR_FAST	

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

	▼
MULTI_USER;	
ONLINE;	
OPEN;	
TRUSTWORTHY;	

GO

Explanation:

USE master;

ALTER DATABASE [DB1] SET

GO

OFFLINE

ONLINE

SINGLE_USER

TRUSTWORTHY

WITH ROLLBACK IMMEDIATE

DBCC CHECKDB ('DB1',

GO

MOINDEX

PHYSICAL_ONLY

REPAIR_ALLOW_DATA_LOSS

REPAIR_FAST

WITH NO_INFOMSGS;

ALTER DATABASE [DB1] SET

GO

MULTI_USER;

ONLINE;

OPEN;

TRUSTWORTHY;

Box 1: SINGLE_USER

The specified database must be in single-user mode to use one of the following repair options.

Box 2: REPAIR_ALLOW_DATA_LOSS

REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.

Note: The REPAIR_ALLOW_DATA_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR_ALLOW_DATA_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.

Reference:

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

NEW QUESTION: 11

You have an on-premises multi-tier application named App1 that includes a web tier, an application tier, and a Microsoft SQL Server tier. All the tiers run on Hyper-V virtual machines.

Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure.

You need to recommend a solution to fail over the database tier of App1 to Azure. The solution must provide the ability to test failover to Azure without affecting the current environment.

What should you include in the recommendation?

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

Answer: D (LEAVE A REPLY)

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Azure Site Recovery is correct because it is the feature whose normal purpose matches the stated requirement. Azure Site Recovery is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure.

Why A is Wrong: Azure Backup is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: Azure Information Protection is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Windows Server Failover Cluster is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

NEW QUESTION: 12

You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.

You need to implement alerts by using Azure Monitor. The solution must meet the following requirements:

Minimize costs.

Aggregate Intelligent Insights telemetry from each database.

What should you do?

A. From the Diagnostic settings of each database, select Send to Log Analytics.

B. From the Diagnostic settings of each database, select Stream to an event hub.

C. From the Diagnostic settings of SQLMI1. select Send to Log Analytics.

D. From the Diagnostic settings of SQLMI1. select Stream to an event hub.

Answer: (SHOW ANSWER)

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> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why A is Correct: Log Analytics stores logs and metrics in a queryable workspace and supports workbooks, dashboards, alerting, and cross-resource analysis. This fits the case because the requested outcome is: You need to implement alerts by using Azure Monitor. The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: From the Diagnostic settings of each database, select Stream to an event hub. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: Log Analytics stores logs and metrics in a queryable workspace and supports workbooks, dashboards, alerting, and cross-resource analysis. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: From the Diagnostic settings of SQLMI1. select Stream to an event hub. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 13

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> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: a private link is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario asks for: 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. That makes a private link the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: a VPN gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: an ExpressRoute gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

NEW QUESTION: 14

You have an Azure subscription that contains an Azure SQL database named DB1 and a managed identity.

You need to create a user in DB1 that will be used by the managed identity to perform automated tasks. How should you complete the query? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

CREATE USER "AutomationAccount"

ASYMETRIC KEY

EXTERNAL PROVIDER

FOR CERTIFICATE

WITH

Key

LOGIN

OBJECT_ID

= '1133a5a2-4c19-47d7-c2d5-42b28997918f'

EXEC sp_

*, "AutomationAccount"

Answer:



Microsoft

CREATE USER "AutomationAccount"

ASYMETRIC KEY
EXTERNAL PROVIDER
FOR CERTIFICATE

WITH = '1133a5a2-4c19-47d7-c2d5-42b28997918f'

EXEC sp_

LOGIN
OBJECT_ID

*, "AutomationAccount"

Explanation:

Answer Area

CREATE USER "AutomationAccount"

FROM

WITH = '1133a5a2-4c19-47d7-c2d5-42b28997918f'

EXEC sp_addrolemember 'db_owner', "AutomationAccount"

Microsoft

Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a user in DB1 that will be used by the managed identity to perform automated tasks.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 15

Task 6

You need to ensure that any enhancements made to the Query Optimizer through patches are available to db1 and db2 on sql37006895.

Answer:

See the explanation part for the complete Solution.

Explanation:

To ensure that any enhancements made to the Query Optimizer through patches are available to db1 and db2 on sql37006895, you need to enable the query optimizer hotfixes option for each database. This option allows you to use the latest query optimization improvements that are not enabled by default¹. You can enable this option by using the ALTER DATABASE SCOPED CONFIGURATION statement².

Here are the steps to enable the query optimizer hotfixes option for db1 and db2 on sql37006895:

* Connect to sql37006895 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

* Open a new query window and run the following commands for each database:

-- Switch to the database context

USE db1;

GO

-- Enable the query optimizer hotfixes option

ALTER DATABASE SCOPED CONFIGURATION SET QUERY_OPTIMIZER_HOTFIXES = ON;

GO

* Repeat the same commands for db2, replacing db1 with db2 in the USE statement.

* To verify that the query optimizer hotfixes option is enabled for each database, you can query the sys.database_scoped_configurations catalog view. The value of the query_optimizer_hotfixes column should be 1 for both databases.

These are the steps to enable the query optimizer hotfixes option for db1 and db2 on sql37006895.

NEW QUESTION: 16

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates.

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

Microsoft

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
- DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
-MaintenanceWindowDuration 120 -PatchCategory "Important"

Get-AzVM -ResourceGroupName "RG1" -Name "SQL" |
- AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
- SQLManagementType
```

- Lightweight

- Full

- Lightweight

- NoAgent

Set-AzVMExtension

Get-AzVMSQLServerExtension

Set-AzVMExtension

Set-AzVMSqlServerExtension

Answer:

Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"  
  
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |  
  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType
```

Set-AzVMExtension
Get-AzVMSQLServerExtension
Set-AzVMExtension
Set-AzVMSqlServerExtension

-Lightweight
-Full
-Lightweight
-NoAgent



Explanation:

Answer Area



```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"  
  
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |  
  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType
```

Set-AzVMExtension

-Lightweight

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

You have an Azure subscription.

You need to deploy two virtual machines named VM1 and VM2 that will host Microsoft SQL Server instances. The solution must meet the following requirements:

- * VM1 must provide the highest I/O throughput-to-vCore ratio.
- * VM2 must provide the highest memory-to-vCore ratio.

Which series should you select for each virtual machine? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

 Microsoft

VM1:

DCadsv5-series
Ebdsv-series
Mv2-series

VM2:

DCadsv5-series
Ebdsv5-series
Mv2-series

Answer:

Answer Area

 Microsoft

VM1:

DCadsv5-series
Ebdsv-series
Mv2-series

VM2:

DCadsv5-series
Ebdsv5-series
Mv2-series

Explanation:

Answer Area

 Microsoft

VM1:

Ebdsv5-series

VM2:

Mv2-series

Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to deploy two virtual machines named VM1 and VM2 that will host Microsoft SQL Server instances.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 18

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.

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

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

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

A. 2

B. 4

C. 8

D. 64

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why C is Correct: 8 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to reduce tempdb contention without negatively affecting server performance. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: 2 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: 4 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: 64 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 19

You have an Azure SQL database named DB1 that contains a nonclustered index named index1.

End users report slow queries when they use index1.

You need to identify the operations that are being performed on the index.

Which dynamic management view should you use?

A.	sys.dm_exec_query_plan_stats
B.	Sys.dm_db_index_physical_stats
C.	Sys.dm_db_index_operational_stats
D.	Sys.dm_db_index_usage_stats

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

Answer: D ([LEAVE A REPLY](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: Option D is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to identify the operations that are being performed on the index.

The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Option A is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Option B is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Option C is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 20

You have an Azure SQL database named SQL1.

You need to implement a disaster recovery solution for SQL1. The solution must minimize the following:

- * The recovery point objective (RPO)
- * The recovery time objective (RTO)
- * Administrative effort

What should you include in the solution?

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

Answer: ([SHOW ANSWER](#)**)**

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why A is Correct: auto failover groups is correct because it is the feature whose normal purpose matches the stated requirement. Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. The scenario wording points to that specific behavior: You need to implement a disaster recovery solution for SQL1.

Why B is Wrong: Azure Site Recovery is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: availability groups is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: active geo - replication is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 21

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: (SHOW ANSWER)

Reference:

https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-ina-data-warehouse Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B and D is Correct: Create a date dimension table that has an integer key in the format of YYYYMMDD. matches the expected DP-300 administration action. Create a date dimension table that has an integer key in the format of YYYYMMDD. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation.

The question is not asking for a general Azure capability; it is asking for the feature that produces this result:

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. Use integer columns for the date fields. matches the expected DP-300 administration action. Use integer columns for the date fields. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: 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.

Why A is Wrong: Create a date dimension table that has a DateTime key. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: Use built-in SQL functions to extract date attributes. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why E is Wrong: Use DateTime columns for the date fields. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 22

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: ([SHOW ANSWER](#))

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why A is Correct: Yes matches the expected DP-300 administration action. Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Each question in the series contains a unique solution that might meet the stated goals.

Why B is Wrong: No is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 23

Task 11

You have a legacy application written for Microsoft SQL Server 2012. The application will be the only application that accesses db1 You need to ensure that db1 is compatible with all the features and syntax of SQL Server 2012.

Answer:

See the explanation part for the complete Solution

Explanation:

To ensure that db1 is compatible with all the features and syntax of SQL Server 2012, you need to set the compatibility level of the database to 110, which is the compatibility level for SQL Server 20121. The compatibility level affects the behavior of certain Transact-SQL statements and features, and determines how the database engine interprets the SQL code2.

You can set the compatibility level of db1 by using the Azure portal or Transact-SQL statements. Here are the steps for both methods:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

Select the database db1 and click on Query Performance Insight in the left menu.

Click on Configure Query Store and select 110 from the Compatibility level dropdown list.

Click on Save to apply the change.

Using Transact-SQL statements:

Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

Open a new query window and run the following command: ALTER DATABASE db1 SET COMPATIBILITY_LEVEL = 110; GO This command will set the compatibility level of db1 to 110, which is equivalent to SQL Server 2012.

These are the steps to set the compatibility level of db1 to 110.

NEW QUESTION: 24

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

What should you run as part of the implementation?

- A.** CREATE LOGIN and the FROM WINDOWS clause
- B.** CREATE USER and the FROM CERTIFICATE clause
- C.** CREATE USER and the FROM LOGIN clause
- D.** CREATE USER and the ASYMMETRIC KEY clause
- E.** CREATE USER and the FROM EXTERNAL PROVIDER clause

Answer: ([SHOW ANSWER](#))

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.) Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user:

```
CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
```

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql> Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why E is Correct: CREATE USER ... FROM EXTERNAL PROVIDER creates contained users mapped to Microsoft Entra identities for Azure SQL authentication. The scenario asks for: You need to implement authentication for ResearchDB1. That makes CREATE USER and the FROM EXTERNAL PROVIDER clause the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: CREATE LOGIN and the FROM WINDOWS clause is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: Certificates support encryption and authentication, but the correct certificate-related object depends on whether the task is TLS, Always Encrypted, or user authentication. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to implement authentication for ResearchDB1.

Why C is Wrong: CREATE USER and the FROM LOGIN clause is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why D is Wrong: CREATE USER and the ASYMMETRIC KEY clause is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

NEW QUESTION: 25

You need to recommend a solution to ensure that the performance of DB3 is optimized after the migration to Azure SQL Database. The solution must meet availability requirements.

What should you include in the recommendation?

- A.** Resource Governor
- B.** a custom resource pool
- C.** vertical scaling
- D.** horizontal scaling

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: vertical scaling is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to recommend a solution to ensure that the performance of DB3 is optimized after the migration to Azure SQL Database. vertical scaling satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Resource Governor is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: a custom resource pool is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: horizontal scaling is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 26

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.

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

Answer:

Deployment option:

	▼
Azure SQL Managed Instance	
SQL Server on Azure Virtual Machines	
An Azure SQL Database single database	

Resiliency option:

	▼
Auto-failover group	
Active geo-replication	
Zone-redundant deployment	

Explanation:
Box 1: SQL Server on Azure VMs
SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.

Box 2: Active geo-replication

Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.

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: 27

You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.

What should you include in the plan?

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

Answer: (SHOW ANSWER)

Reference:

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

Topic 2, Contoso LtdCase study

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.

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

Existing Environment

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.

Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

Name	Role
SQL1	Primary data warehouse
SQL2	Secondary data warehouse
SQL3	Extract, transform, and load (ETL) server

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.

The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.

Current Business Model

Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers. The customers receive the files by using FTP.

Requirements

Planned Changes

Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.

Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.

Start onboarding customers to the new PaaS solution within six months.

Business Goals

Contoso identifies the following business requirements:

Use built-in Azure features whenever possible.

Minimize development effort whenever possible.

Minimize the compute costs of the PaaS solutions.

Provide all the customers with their own copy of the database by using the PaaS solution.

Provide the customers with different table and row access based on the customer's service agreement.

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.

Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso.

Technical Requirements

Contoso identifies the following technical requirements:

Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.

All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion.

Users must be able to review the queries issued against the PaaS databases and identify any new objects created.

Downtime during the PostgreSQL database migration must be minimized.

Monitoring Requirements

Contoso identifies the following monitoring requirements:

Notify administrators when a PaaS database has a higher than average CPU usage.

Use a single dashboard to review security and audit data for all the PaaS databases.

Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.

Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.


PaaS Prototype

During prototyping of the PaaS solution in Azure, you record the compute utilization of a customer's Azure SQL database as shown in the following exhibit.

Show data for last: 1 hour 24 hours **7 days**

Aggregation type: Max ▾

Compute utilization



Role Assignments

For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

[+ Add](#)
[Edit columns](#)
[Refresh](#)
[Remove](#)
[Got feedback?](#)

[Check access](#)
[Role assignments](#)
[Deny assignments](#)
[Classic administrators](#)
[Roles](#)

Manage access to Azure resources for users, groups, service principals and managed identities at this scope by creating role assignments. [Learn more](#)

Number of role assignments for this subscription 15 2000

Name
 Type
 Role
 Scope

Group by

Showing a filtered set of results. Total number of role assignments: 15

2 items (2 Groups)

☐	Name	Type	Role	Scope
☐	 DBAGroup1	Group	Contributor	This resource
☐	 DBAGroup2	Group	SQL DB Contributor	This resource

NEW QUESTION: 28

You have an Azure subscription that contains an Azure SQL database. The database contains a table named table1 that uses partitioned columnstores. You need to configure table1 to meet the following requirements:

- * Each partition must be compressed
- * The compression ratio must be maximized.
- * You must be able to index the compressed data.

What should you use?

- A.** page compression
- B.** columnstore compression
- C.** columnstore archival compresssion
- D.** GZIP compression

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why C is Correct: columnstore archival compresssion is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to configure table1 to meet the following requirements: * Each partition must be compressed * The compression ratio must be maximized.

The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: page compression is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why B is Wrong: columnstore compression is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: GZIP compression is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 29

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.** sp_addremotelogin
- B.** xp_instance_regwrite
- C.** sp_cnarge_users_login
- D.** xp_grant_login

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why B is Correct: xp_instance_regwrite matches the expected DP-300 administration action.

xp_instance_regwrite is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to configure SQL! to use mixed mode authentication.

Why A is Wrong: sp_addremotelogin is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why C is Wrong: sp_cnarge_users_login is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: xp_grant_login is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 30

Vou have an Azure SQL database named DB1.

You have 10 Azure virtual machines that connect to a virtual network subnet named Subnet 1.

You need to implement a database-level firewall that meets the following requirements:

- * Ensures that only the 10 virtual machines can access DB1

* Follows the principle of least privilege
How should you configure the firewall rule, and how should you establish network connectivity from the virtual machines to DB1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

Firewall rule:

Add your client IPv4 address.
Add your client IPv4 address.
Allow Azure services and resources to access DB1.
Allow traffic from a specific virtual network.

Network connectivity:

Create a private endpoint.
Assign static public IP addresses to the virtual machines.
Create a private endpoint.
Create a service endpoint.

Answer:

Answer Area

Microsoft

Firewall rule:

Add your client IPv4 address.
Add your client IPv4 address.
Allow Azure services and resources to access DB1.
Allow traffic from a specific virtual network.

Network connectivity:

Create a private endpoint.
Assign static public IP addresses to the virtual machines.
Create a private endpoint.
Create a service endpoint.

Explanation:

Answer Area

Microsoft

Firewall rule:

Add your client IPv4 address.

Network connectivity:

Create a private endpoint.

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.
Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to implement a database-level firewall that meets the following requirements: * Ensures that only the 10 virtual machines can access DB1 * Follows the principle of least privilege
How should you configure the firewall rule, and how should you establish network connectivity from the virtual machines to DB1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 31

You configure backup for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore
Specify how long you want to keep your point-in-time backups. [Learn more](#)

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

Long-term retention
Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

Weekly LTR Backups
Keep weekly backups for:

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

Microsoft LTR Backups
Keep an annual backup for:

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

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

Answer Area

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

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

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

65 backup copies
1 backup copy
52 backup copies
64 backup copies
65 backup copies

Answer:

Answer Area

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

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

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

65 backup copies
1 backup copy
52 backup copies
64 backup copies
65 backup copies

Explanation:

Answer Area

To restore from a failure that occurred two days ago and caused minimal data loss, you must use a [answer choice] point-time restore (PITR) backup.

After the 52nd weekly backup runs, there will be [answer choice] in long term retention. 65 backup copies

Microsoft

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You configure backup for an Azure SQL database as shown in the following exhibit.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 32

You plan to deploy three instances of SQL Server on Azure Virtual Machines that will each contain 20 databases. You need to recommend a solution that meets the following requirements:

- * Ensures that the deployment is highly available
 - * Minimizes administrative effort to manage users, logins, permissions, and SQL Server Agent jobs across the instances
- What should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Type of availability group:

Minimum Microsoft SQL Server version:

Microsoft

Answer:

Answer Area

Type of availability group:
Contained
Distributed

Minimum Microsoft SQL Server version:
2017
2019
2022

Microsoft

Explanation:

Answer Area

Type of availability group:

Minimum Microsoft SQL Server version:


Microsoft

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to recommend a solution that meets the following requirements: * Ensures that the deployment is highly available * Minimizes administrative effort to manage users, logins, permissions, and SQL Server Agent jobs across the instances What should you inc Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 33

You have a Microsoft SQL Server database named DB1 that contains a table named Table1.
The database role membership for a user named User1 is 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.

User1 can [answer choice].

add a column to Table1
delete a row from Table1
delete Table1

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.



db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

Answer:

User1 can [answer choice].

add a column to Table1
delete a row from Table1
delete Table1



To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

Explanation:



User1 can [answer choice].

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

add a column to Table1
delete a row from Table1
delete Table1

db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

Box 1: delete a row from Table1

Members of the db_datawriter fixed database role can add, delete, or change data in all user tables.

Box 2: db_datareader

Members of the db_datareader fixed database role can read all data from all user tables.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles> Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have a Microsoft SQL Server database named DB1 that contains a table named Table1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 34

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: (SHOW ANSWER)

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>

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why A is Correct: automated can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. In this scenario, the important constraint is: Which type of Databricks cluster should you use? automated satisfies that constraint without adding an unrelated service or manual process.

Why B is Wrong: interactive can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why C is Wrong: High Concurrency can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

NEW QUESTION: 35

You have an on-premises Microsoft SQL Server instance named SQL1 that hosts multiple databases. You have an Azure subscription that contains an Azure SQL managed instance named SQLM11. You need to perform an offline migration of the SQLi databases to SQLMH by using Azure Data Studio. What should you install on SQL1?

- A. Data Migration Assistant (DMA)
- B. a self-hosted integration runtime
- C. SQL Server Data Tools (SSDT)
- D. an Azure Migrate appliance

Answer: (SHOW ANSWER)

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: a self-hosted integration runtime is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario asks for: You need to perform an offline migration of the SQLi databases to SQLMH by using Azure Data Studio. That makes a self-hosted integration runtime the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Data Migration Assistant (DMA) is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: SQL Server Data Tools (SSDT) is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: an Azure Migrate appliance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

NEW QUESTION: 36

You have an Azure subscription that contains a user named user1 and an Azure SQL database named OBI.

You configure long-term retention (LTR) for DB1.

You need to ensure that User1 can delete LTR backups of DB1. The solution must follow the principle of least privilege. Which role should you assign to User1?

A. SQL Server Contributor

B. Owner

C. SQL Security Manager

D. Contributor

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why B is Correct: Owner is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario asks for: You need to ensure that User1 can delete LTR backups of DB1. That makes Owner the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: SQL Server Contributor is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that User1 can delete LTR backups of DB1.

Why C is Wrong: SQL Security Manager is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: Contributor is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that User1 can delete LTR backups of DB1.

NEW QUESTION: 37

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.

ACTIONS

Duplicate Azure Policy definitions

Run Azure Policy remediation tasks

Create an Azure Blueprints assignment

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Answer Area

>

<

Microsoft

<

>

Answer:

Actions

Duplicate Azure Policy definitions

Run Azure Policy remediation tasks

Create an Azure Blueprints assignment

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Answer Area

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Run Azure Policy remediation tasks

Explanation:

Create an Azure Policy initiative

Create an Azure Policy initiative assignment

Run Azure Policy remediation tasks

Step 1: Create an Azure Policy Initiative

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take. With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

Step 2: Create an Azure Policy Initiative assignment

Assign the initiative definition you created in the previous step.

Step 3: Run Azure Policy remediation tasks

To apply the Policy Initiative to the existing SQL databases.

Reference:

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

NEW QUESTION: 38

You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1. You need to ensure that app1 can access db1. The solution must use the principle of least privilege. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

CREATE

USER

CREDENTIAL

LOGIN

USER

[app1]

FROM EXTERNAL PROVIDER

FOR LOGIN app1

FROM EXTERNAL PROVIDER

FROM LOGIN app1

WITHOUT LOGIN

Answer:

Answer Area

CREATE

USER

CREDENTIAL

LOGIN

USER

[app1]

FROM EXTERNAL PROVIDER

FOR LOGIN app1

FROM EXTERNAL PROVIDER

FROM LOGIN app1

WITHOUT LOGIN

Explanation:

Answer Area

CREATE

USER

[app1]

FROM EXTERNAL PROVIDER

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles. Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that app1 can access db1. Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

NEW QUESTION: 39

You configure version control for an Azure Data Factory instance as shown in the following exhibit.

Home

Connections

Linked services

Integration runtimes

Azure Purview (Preview)

Source control

Git configuration

ARM template

Parameterization template

Author

Triggers

Global parameters

Security

Customer managed key

Managed private endpoints

Git repository

Git repository information associated with your data factory. [CI/CD best practices](#)

Setting

Disconnect

Repository type

Azure DevOps Git

Azure DevOps Account

CONTOSO

Project name

Data

Repository name

dwh_batchetl

Collaboration branch

main

Publish branch

adf_publish

Root folder

/

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 for 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:

Azure Resource Manager (ARM) templates for 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

Explanation:



Box 1: adf_publish

By default, data factory generates the Resource Manager templates of the published factory and saves them into a branch called adf_publish. To configure a custom publish branch, add a publish_config.json file to the root folder in the collaboration branch. When publishing, ADF reads this file, looks for the field publishBranch, and saves all Resource Manager templates to the specified location. If the branch doesn ' t exist, data factory will automatically create it. And example of what this file looks like is below:

```
{
"publishBranch " : " factory/adf_publish "
}
```

Box 2: /dwh_barchlet/ adf_publish/contososales

RepositoryName: Your Azure Repos code repository name. Azure Repos projects contain Git repositories to manage your source code as your project grows. You can create a new repository or use an existing repository that ' s already in your project.

Reference:

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

NEW QUESTION: 40

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

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

What should you include in the recommendation?

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: transactional replication to an Azure SQL Managed Instance matches the expected DP-300 administration action. Transactional replication distributes committed changes to subscribers, but it is not the default full-database migration method for Azure SQL Managed Instance when the requirement is minimal effort. The question is not asking for a general Azure capability; it is asking for the feature that produces this result:

You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database.

Why A is Wrong: database mirroring on an instance of SQL Server on Azure Virtual Machines is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: availability groups for SQL Server on Azure Virtual Machines is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: an Azure SQL Managed Instance link is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 41

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. Decrease the auto-pause delay for db1.
 - C. Disable auto-pause delay for db1.
 - D. Enable automatic tuning for db1.

Answer: (SHOW ANSWER)

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. This fits the case because the requested outcome is: You need to modify db1 to resolve the issues experienced by Appl as soon as possible, without considering immediate costs What you do? The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why B is Wrong: Decrease the auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: Disable auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 42

You have a Microsoft SQL Server 2022 database named DB1 that contains a table named Customers. Customers contains the following columns.

Name	Data type	Data sample
Email	nvarchar(200)	jane.doe@domain.com john.doe.@domain.com eric.doe@domain.com
CreditcardNumber	char(19)	4444 5555 6666 7777 9999 8888 7777 6666 1114 3746 2845 2734
Birthdate	date	1979-03-16 1987-01-20 2001-10-27

You need to enable dynamic data masking for Customers. The solution must provide an output for each column as shown in the following table.

Column	Mask
Email	First letter, first letter in domain, and domain visible
CreditcardNumber	All digits masked
Birthdate	Year and month visible

The solution must minimize development effort.

Which masking function should you use for each column? To answer, drag the appropriate masking functions to the correct columns. Each masking function 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.

Masking functions

Custom string

Datetime

Default

Email

Random

Answer Area

Email:

CreditcardNumber:

Birthdate:

Microsoft

Answer:

Masking functions

Custom string

Datetime

Default

Email

Random

Answer Area

Email:

CreditcardNumber:

Birthdate:

Microsoft

Explanation:

Masking functions

Custom string

Datetime

Default

Email

Random

Answer Area

Email:

CreditcardNumber:

Birthdate:

Microsoft

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to enable dynamic data masking for Customers.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 43

You have an Azure subscription that uses a domain named contoso.com.

You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.

DBServer1 has a high volume of data changes and low latency requirements for data writes.

You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.

What should you do?

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

Answer: C ([LEAVE A REPLY](#))

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-ver15> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. This fits the case because the requested outcome is: DBServer1 has a high volume of data changes and low latency requirements for data writes. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why B is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

NEW QUESTION: 44

You need to recommend a backup solution to restore DB3. The solution must meet the availability requirements. Which type of backup should you use?

- A.** transaction log
- B.** point-in-time restore (PITR)
- C.** differential
- D.** long-term retention (LTR)

Answer: ([SHOW ANSWER](#)**)**

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: differential is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario asks for: You need to recommend a backup solution to restore DB3. That makes differential the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: transaction log is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: point-in-time restore (PITR) is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: long-term retention (LTR) is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 45

You have an Azure subscription.

You need to create a new database named DB1 in the subscription. The solution must meet the following requirements:

Ensure that tables created in DB1 automatically provide cryptographic proof when data is altered.

Ensure that auditors can verify whether the data has NOT been modified.

Which T-SQL command should you run?

A. CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256 ENCRYPTION BY SERVER CERTIFICATE

B. CREATE DATABASE DB1 WITH TRUSTWORTHY

C. CREATE DATABASE DB1 WITH LEDGER = ON

D. CREATE DATABASE DB1 WITH EDITION = ' GeneralPurpose ' , SERVICE_OBJECTIVE = ' GP_Gen5_2 '

Answer: (SHOW ANSWER)

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why C is Correct: CREATE DATABASE DB1 WITH LEDGER = ON is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to create a new database named DB1 in the subscription. CREATE DATABASE DB1 WITH LEDGER = ON satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Certificates support encryption and authentication, but the correct certificate-related object depends on whether the task is TLS, Always Encrypted, or user authentication. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why B is Wrong: CREATE DATABASE DB1 WITH TRUSTWORTHY is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: CREATE DATABASE DB1 WITH EDITION = ' GeneralPurpose ' , SERVICE_OBJECTIVE = ' GP_Gen5_2 ' is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 46

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: (SHOW ANSWER)

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/database/sql-server-to-sql-database- overview> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: SQL Server on an Azure Virtual Machine is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario.

This fits the case because the requested outcome is: Which service should you use? The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Azure SQL Database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Azure SQL Managed Instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: Azure Database for MySQL is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (**503** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 47

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

- New data is accessed frequently and must be available as quickly as possible.
- Data that is older than five years is accessed infrequently but must be available within one second when requested.
- Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
- Costs must be minimized while maintaining the required availability.

How should you manage the data? To answer, select the appropriate options in the answer area.


NOTE: Each correct selection is worth one point.

The screenshot shows a Microsoft Azure portal interface. It features two sections for data archiving. The first section, labeled 'Five-year-old data:', has a dropdown menu with four options: 'Delete the blob.', 'Move to archive storage.', 'Move to cool storage.', and 'Move to hot storage.'. The second section, labeled 'Seven-year-old data:', also has a dropdown menu with the same four options. A large, semi-transparent watermark reading 'freecram.net' is overlaid diagonally across the center of the image.

Answer:



Explanation:



Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery

Older data not used frequently but expected to be available immediately when accessed Large data sets that need to be stored cost effectively, while more data is being gathered for future processing Note: Hot -

Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Move to archive storage

Example usage scenarios for the archive access tier include:

Long-term backup, secondary backup, and archival datasets

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

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

NEW QUESTION: 48

You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.

You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service. How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Backup type:

Full and log backups only

Full backup only

Log backup only

Backup option:

WITH CHECKSUM

WITH NOINIT

WITH UNLOAD

Answer:

Backup type:

Full and log backups only

Full backup only

Log backup only

Backup option:

WITH CHECKSUM

WITH NOINIT

WITH UNLOAD

Explanation:

Backup type:

Full and log backups only

Full backup only

Log backup only

Backup option:

WITH CHECKSUM

WITH NOINIT

WITH UNLOAD

Box 1: Full and log backups only

Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn ' t support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM

Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

NEW QUESTION: 49

You deploy an instance of SQL Server on Azure Virtual Machines named VM1.

You need to create a SQL Server Agent job that will rebuild indexes of the databases hosted on VM1. The solution must use the principle of least privilege.

What should you create first?

- A. an Elastic Job agent
- B. a local Windows account
- C. a system-assigned managed identity in Microsoft Entra ID
- D. a user-assigned managed identity in Microsoft Entra ID

Answer: (SHOW ANSWER)

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why C is Correct: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. The scenario asks for: You need to create a SQL Server Agent job that will rebuild indexes of the databases hosted on VM1. That makes a system-assigned managed identity in Microsoft Entra ID the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Elastic Jobs automate T-SQL execution across many Azure SQL databases, especially for multi-tenant or shard-style deployments. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why B is Wrong: a local Windows account is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 50

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> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. This fits the case because the requested outcome is: 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. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: a VPN gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: a private link is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: an ExpressRoute gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

NEW QUESTION: 51

You have an Azure subscription.

You need to deploy an Azure SQL managed instance that meets the following requirements:

- * Optimize latency.
- * Maximize the memory-to-vCore ratio.

Which service tier and hardware generation should you use? To answer, select the apocopate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Service tier

Business Critical

General Purpose

Hyperscale

Hardware generation:

Premium-series - memory optimized


Premium-series

Standard-series (Gen 5)

Premium-series - memory optimized

Answer:

Answer Area



Explanation:



Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to deploy an Azure SQL managed instance that meets the following requirements: * Optimize latency.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 52

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: (SHOW ANSWER)

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: 53

You have an Azure subscription.

You need to provision a logical server for Azure SQL Database by using a Bicep template.

how should you complete the template? To answer, select the appropriate options in the answer are NOTE Each correct selection is worth one point.



Microsoft

```
param serverName string = uniqueString('sql', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string

@allowed()
@secure0
@sys

param administratorLoginPassword string

resource sqlServer 'Microsoft.Sql/servers@2024-05-01-preview' = {
  output
  param
  resource

  administratorLogin: administratorLogin
  administratorLoginPassword: administratorLoginPassword
}
```

resource

output

param

resource

Answer:

Answer Area

```
param serverName string = uniqueString('sql', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string
  @allowed()
  @secure()
  @sys
param administratorLoginPassword string
  resource sqlServer 'Microsoft.Sql/servers@2024-05-01-preview' = {
    output
    param
    resource
  }
  administratorLogin: administratorLogin
  administratorLoginPassword: administratorLoginPassword
```



Explanation:

Answer Area

```
param serverName string = uniqueString('sql', resourceGroup().id)
param location string = resourceGroup().location
param administratorLogin string
  @secure()
param administratorLoginPassword string
  resource sqlServer 'Microsoft.Sql/servers@2024-05-01-preview' = {
    name: serverName
    location: location
    properties: {
      administratorLogin: administratorLogin
      administratorLoginPassword: administratorLoginPassword
    }
  }
```

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to provision a logical server for Azure SQL Database by using a Bicep template.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 54

You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.

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

NOTE: Each correct selection is worth one point.

Table type to store retail store data:

Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Answer:

Table type to store retail store data:

Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Explanation:

Table type to store retail store data:

Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Box 1: Hash

Scenario:

Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.

A hash distributed table can deliver the highest query performance for joins and aggregations on large tables.

Box 2: Round-robin

Scenario:

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

A round-robin table is the most straightforward table to create and delivers fast performance when used as a staging table for loads. These are some scenarios where you should choose Round robin distribution:

When you cannot identify a single key to distribute your data.

If your data doesn't frequently join with data from other tables.

When there are no obvious keys to join.

Reference:

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

NEW QUESTION: 55

Task 2

You need to configure differential backups for the db1 Azure SQL database to be once a day instead of twice day. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

See the explanation part for the complete Solution.

Explanation:

Requirement: Configure differential backups for Azure SQL Database db1 to run once a day instead of twice a day.

Correct setting: Change Differential backup frequency from 12 Hours to 24 Hours.

Azure SQL Database supports differential backup frequency of either 12 hours or 24 hours. A 12-hour frequency means twice per day; a 24-hour frequency means once per day. Microsoft also notes that 24-hour differential backup frequency can increase restore time compared with 12-hour frequency.

Method 1 - Azure Portal Method

This is the best method for the simulation.

Step 1: Open the Azure SQL logical server

Sign in to the Azure portal.

Search for SQL servers.

Open the logical SQL server that hosts database db1.

Do not start from SQL Server Management Studio for this task. The differential backup frequency is an Azure SQL backup policy setting, not a normal T-SQL database setting.

Step 2: Open the Backups page

In the SQL server left menu, select Backups.

Select the Retention policies tab.

Microsoft's documented portal path is to go to the logical SQL server, select Backups, then select the Retention policies tab.

Step 3: Select database db1

In the list of databases, locate db1.

Select the checkbox next to db1.

Select Configure policies from the action bar.

Step 4: Change the differential backup frequency

In the policy configuration pane:

Find Differential backup frequency.

Change it from:

12 Hours

to:

24 Hours

Leave the PITR retention period unchanged unless the task specifically tells you to change it.

Select Apply or Save.

Microsoft's documented option is exactly 12 Hours or 24 hours under Differential backup frequency.

Method 2 - Azure CLI Method

Use this if the simulation provides Cloud Shell.

```
az sql db str-policy set \  
--resource-group < resource-group-name > \  
--server < server-name > \  
--name db1 \  
--retention-days < current-retention-days > \  
--diffbackup-hours 24
```

Example:

```
az sql db str-policy set \  
--resource-group RG1 \  
--server sql60152867 \  
--name db1 \  
--retention-days 7 \  
--diffbackup-hours 24
```

Microsoft documents `az sql db str-policy set` with `--diffbackup-hours 24` for changing active database differential backup frequency. Valid values are 12 or 24 hours.

Be careful: do not guess the retention days blindly in a real environment. In the exam lab, use the existing retention value shown in the portal unless the task also asks you to change retention.

Method 3 - PowerShell Method

Use this if Azure PowerShell is available.

```
Set-AzSqlDatabaseBackupShortTermRetentionPolicy `  
-ResourceGroupName " < resource-group-name > " `  
-ServerName " < server-name > " `  
-DatabaseName " db1 " `  
-RetentionDays < current-retention-days > `  
-DiffBackupIntervallnHours 24
```

Microsoft documents `Set-AzSqlDatabaseBackupShortTermRetentionPolicy` with `-DiffBackupIntervallnHours 24` for setting Azure SQL Database differential backup frequency.

SSMS / T-SQL Clarification

For this task, SSMS is not the right tool to change the setting.

There is no normal ALTER DATABASE T-SQL command in Azure SQL Database to change automated differential backup frequency. Microsoft documents this change through:

Azure portal

Azure CLI

PowerShell

REST API

not SSMS/T-SQL.

You may use SSMS only to confirm the database exists and is accessible, but the backup frequency setting must be changed from Azure management tools.

NEW QUESTION: 56

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.

EXECUTE AS

▼

= 'user1@contoso.com'

CALLER

LOGIN

OWNER

USER

GO

SELECT SUSER_SNAME()

▼

REVERT

REVOKE

ROLLBACK

GO

 Microsoft

Answer:

EXECUTE AS

▼

= 'user1@contoso.com'

CALLER

LOGIN

OWNER

USER

GO

SELECT SUSER_SNAME()

▼

REVERT


REVOKE

ROLLBACK

GO

 Microsoft

Explanation:



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> Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to test impersonation of user1 in db1 by running a SELECT statement and returning to the original execution context.

Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

NEW QUESTION: 57

You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2.

Which statement should you run?

- A. BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net/mycontainer /DB1.bak ' with (Differential, STATS=5, COMPRESSION);
- B. BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net/mycontainer /DB1.bak ' with (COPY_ONLY, STATS=5, COMPRESSION);
- C. BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net/mycontainer /DB1.bak ' with (File_Snapshot, STATS=5, COMPRESSION);
- D. BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net/mycontainer /DB1.bak ' with (NoInit, STATS=5, COMPRESSION);

Answer: (SHOW ANSWER)

BACKUP DATABASE supports only copy-only full backups of databases, files, or filegroups when it ' s executed on secondary replicas. Copy-only backups don ' t impact the log chain or clear the differential bitmap.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why B is Correct: BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net /mycontainer/DB1.bak ' with (COPY_ONLY, STATS=5, COMPRESSION); is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. This fits the case because the requested outcome is: You need to initiate a full backup of DB1 on SQL2. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net /mycontainer/DB1.bak ' with (Differential, STATS=5, COMPRESSION); is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net /mycontainer/DB1.bak ' with (File_Snapshot, STATS=5, COMPRESSION); is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: BACKUP DATABASE DB1 TO URL= ' https://mystorageaccount.blob.core.windows.net /mycontainer/DB1.bak ' with (NoInit, STATS=5, COMPRESSION); is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

NEW QUESTION: 58

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/> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: Live Query Statistics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to identify which operation in the query is causing the performance issue. Live Query Statistics satisfies that constraint without adding an unrelated service or manual process.

Why B is Wrong: an estimated execution plan is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: an actual execution plan is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Client Statistics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 59

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> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: SNAPSHOT matches the expected DP-300 administration action. SNAPSHOT is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which isolation level should you configure?

Why A is Wrong: SERIALIZABLE is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: READ UNCOMMITTED is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: READ COMMITTED is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

NEW QUESTION: 60

You have an Azure SQL managed instance named sqldbmi1 that contains a database name Sales.
You need to initiate a backup of Sales.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

BACKUP DATABASE Sales

Microsoft

TO DISK = '\\BackupSystem\BackupDisk1\Sales.bak'

TO DISK = 'X:\BAK\Sales.bak'

TO 'Sales_Backup'

TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

COPY_ONLY;

ENCRYPTION;

FILE_SNAPSHOT;

NO_TRUNCATE;

Answer:

Microsoft

BACKUP DATABASE Sales

TO DISK = '\\BackupSystem\BackupDisk1\Sales.bak'

TO DISK = 'X:\BAK\Sales.bak'

TO 'Sales_Backup'

TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

COPY_ONLY;

ENCRIPTION;

FILE_SNAPSHOT;

NO_TRUNCATE;

Explanation:

Microsoft

BACKUP DATABASE Sales

TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'

WITH STATS = 5,

COPY_ONLY;

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to initiate a backup of Sales.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 61

You have 20 Azure SQL databases provisioned by using the vCore purchasing model.

You plan to create an Azure SQL Database elastic pool and add the 20 databases.

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

NOTE: Each correct selection is worth one point.

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

Answer: (SHOW ANSWER)

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

For vCore-based purchasing model: $\text{MAX}(\text{Total number of DBs} \times \text{average vCore utilization per DB}, \text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB})$ A: Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why A and C and E is Correct: total size of all the databases matches the expected DP-300 administration action. total size of all the databases is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which three metrics should you use to size the elastic pool to meet the demands of your workload? number of concurrently peaking databases * peak CPU utilization per database matches the expected DP-300 administration action. CPU percentage is important for database compute monitoring, but it is not the Stream Analytics sizing metric when the question is about streaming units. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which three metrics should you use to size the elastic pool to meet the demands of your workload? total number of databases * average CPU utilization per database is correct because it is the feature whose normal purpose matches the stated requirement. CPU percentage is important for database compute monitoring, but it is not the Stream Analytics sizing metric when the question is about streaming units. The scenario wording points to that specific behavior: Which three metrics should you use to size the elastic pool to meet the demands of your workload?

Why B is Wrong: Active geo-replication creates readable secondary databases, but failover groups add the managed listener and automatic failover abstraction needed for application continuity at scale. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why D is Wrong: A session window groups events separated by periods of inactivity, which is useful for user sessions rather than fixed continuous aggregation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (**503** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 62

You have an on-premises server named Server1 that has Microsoft SQL Server installed. Server1 contains a database named DB1. You have an Azure subscription that contains a user named User1. You need to ensure that User1 can perform the following actions:

- * Register a resource provider to support migrating DB1 to Azure.
- * Create an Azure Database Migration Service instance
- * Migrate DB1 to Azure SQL Database.

The solution must follow the principle of least privilege.

Which role should you assign to User1 for the subscription, and which resource provider should User1 register? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

Role:

- Contributor
- SQL DB Contributor
- SQL Server Contributor
- SQL Security Manager

Resource provider:

- Microsoft.Databox
- Microsoft.DataMigration
- Microsoft.Migrate

Answer:

Answer Area

Microsoft

Role:

- Contributor
- SQL DB Contributor
- SQL Server Contributor
- SQL Security Manager

Resource provider:

- Microsoft.Databox
- Microsoft.DataMigration
- Microsoft.Migrate

Explanation:

Answer Area

Microsoft

Role: Contributor

Resource provider: Microsoft.DataMigration

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that User1 can perform the following actions: * Register a resource provider to support migrating DB1 to Azure.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

NEW QUESTION: 63

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

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

You have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Must use an Azure Data Factory, not an Azure Databricks job.

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why B is Correct: No can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. This fits the case because the requested outcome is: Each question in the series contains a unique solution that might meet the stated goals. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Yes can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

NEW QUESTION: 64

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: ([SHOW ANSWER](#)**)**

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: 65

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

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

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

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

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

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

a header
a data
an empty

Answer:


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

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

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

a header
a data
an empty

Explanation:



Box 1: CSV files that have file named beginning with " tripdata_2020 "

Box 2: a header

FIRSTROW = ' first_row '

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 66

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: (SHOW ANSWER)

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> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why D is Correct: The serverless tier automatically scales compute and can pause during inactive periods, making it suitable for intermittent workloads with bursts and idle windows. The scenario asks for: Based on the PaaS prototype, which Azure SQL Database compute tier should you use? That makes Serverless the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.

Why B is Wrong: Hyperscale is designed for very large databases and fast scale-out storage architecture, not merely for intermittent CPU bursts. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.

Why C is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It applies to a different Azure data-platform design than the one the question is testing.

Verification Note: The original listed answer was A; the corrected answer is D because The prototype shows intermittent spikes with idle periods; the serverless compute tier is designed for intermittent usage and automatic pause/resume cost savings.

NEW QUESTION: 67

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

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

You have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why A is Correct: Yes matches the expected DP-300 administration action. Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait /blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Each question in the series contains a unique solution that might meet the stated goals.

Why B is Wrong: No is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 68

You are building a database backup solution for a SQL Server database hosted on an Azure virtual machine.

In the event of an Azure regional outage, you need to be able to restore the database backups. The solution must minimize costs.

Which type of storage accounts should you use for the backups?

A. locally-redundant storage (LRS)

B. read-access geo-redundant storage (RA-GRS)

C. zone-redundant storage (ZRS)

D. geo-redundant storage

Answer: ([SHOW ANSWER](#))

Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages. However, that data is available to be read only if the customer or Microsoft initiates a failover from the primary to secondary region. When you enable read access to the secondary region, your data is available to be read if the primary region becomes unavailable. For read access to the secondary region, enable read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Reference:

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

NEW QUESTION: 69

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1. You need to configure SQLMI1 to use the Business Critical service tier Which PowerShell cmdlet should you run?

- A. Set-AzSqlServerConfigurationOption
- B. Update-SqlvM
- C. Set-AiSqlDatabase
- D. Set-AzSqlInstance

Answer: D (LEAVE A REPLY)

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why D is Correct: Set-AzSqlInstance matches the expected DP-300 administration action. Set-AzSqlInstance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to configure SQLMI1 to use the Business Critical service tier Which PowerShell cmdlet should you run?

Why A is Wrong: Set-AzSqlServerConfigurationOption is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: Update-SqlvM is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why C is Wrong: Set-AiSqlDatabase is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

NEW QUESTION: 70

You have an Azure subscription that contains an Always On availability group named AG1. AG1 contains three replicas named Server1. Server2. and Server3.

You execute the following T-SQL statements on Server1.

```
SELECT
Groups.[Name] AS 'Availability group',
Cluster.database_name AS DBName
FROM sys.dm_hadr_availability_group_states States
INNER JOIN master.sys.availability_groups Groups ON States.group_id = Groups.group_id
INNER JOIN sys.availability_databases_cluster Cluster ON Groups.group_id = Cluster.group_id
WHERE primary_replica != @@Servername
ORDER BY
Availability ASC,
DBName ASC;
```

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The statement will only return databases that are members of AG1.	☐	☐
The statement will return all the databases for which Server1 is the secondary replica.	☐	☐
The statement will return all the Always On availability groups for which Server3 is a secondary replica.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The statement will only return databases that are members of AG1.	☒	☐
The statement will return all the databases for which Server1 is the secondary replica.	☒	☐
The statement will return all the Always On availability groups for which Server3 is a secondary replica.	☐	☒

Explanation:

Answer Area

Statements	Yes	No
The statement will only return databases that are members of AG1.	☒	☐
The statement will return all the databases for which Server1 is the secondary replica.	☒	☐
The statement will return all the Always On availability groups for which Server3 is a secondary replica.	☐	☒

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure subscription that contains an Always On availability group named AG1.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

NEW QUESTION: 71

You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance requirements.

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

Actions

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

Answer Area

>

<

↑

↓

Answer:

Actions

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

Answer Area

Create an Azure AD administrator for the logical server

Create contained database users

Connect to the databases by using an Azure AD account

Explanation:

Create an Azure AD administrator for the logical server

Create contained database users

Connect to the databases by using an Azure AD account

Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

Create and populate Azure AD.

Optional: Associate or change the active directory that is currently associated with your Azure Subscription.

Create an Azure Active Directory administrator. (Step 1)

Configure your client computers.

Create contained database users in your database mapped to Azure AD identities. (Step 2) Connect to your database by using Azure AD identities. (Step 3) Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview>

NEW QUESTION: 72

You have an Azure SQL managed instance.

You need to restore a database named DB1 by using Transact-SQL.

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

NOTE: Each correct selection is worth one point.

RESTORE

▼

 DB1 FROM

▼

DATABASE

FILE

LOG

DISK = N'\\NAS01\$SQLBackups\DB1.bak';

TAPE = N'\\Tape0'

URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Answer:

RESTORE

▼

 DB1 FROM

▼

DATABASE

FILE

LOG

DISK = N'\\NAS01\$SQLBackups\DB1.bak';

TAPE = N'\\Tape0'

URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Explanation:

RESTORE

▼

 DB1 FROM

▼

DATABASE

FILE

LOG

DISK = N'\\NAS01\$SQLBackups\DB1.bak';

TAPE = N'\\Tape0'

URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to restore a database named DB1 by using Transact-SQL.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 73

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: ([SHOW ANSWER](#))

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/>

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why A is Correct: an annotation matches the expected DP-300 administration action. an annotation is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to label each pipeline with its main purpose of either ingest, transform, or load.

Why B is Wrong: a resource tag is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: a run group ID is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: a user property is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why E is Wrong: a correlation ID is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 74

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/>

NEW QUESTION: 75

You have an Azure SQL database named DB1 in the General Purpose service tier.

The performance metrics for DB1 are shown in the following exhibit.



You need to reduce the Log 10 percentage. The solution must minimize costs.

What should you do?

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D is Correct: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. This fits the case because the requested outcome is: You need to reduce the Log 10 percentage. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It applies to a different Azure data-platform design than the one the question is testing.

Why B is Wrong: B. Change Recovery model to Simple. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Perform a checkpoint operation. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

NEW QUESTION: 76

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.

You need to ensure that users from each company can view only the data of their respective company.

Which two objects should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a column encryption key
- B. asymmetric keys
- C. a function
- D. a custom role-based access control (RBAC) role
- E. a security policy

Answer: (SHOW ANSWER)

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes.

Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac>

<https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D and E is Correct: a custom role-based access control (RBAC) role is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. This fits the case because the requested outcome is: You need to ensure that users from each company can view only the data of their respective company. The selected feature addresses that outcome directly rather than relying on a workaround. a security policy is correct because it is the feature whose normal purpose matches the stated requirement. a security policy is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: You need to ensure that users from each company can view only the data of their respective company.

Why A is Wrong: a column encryption key is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: asymmetric keys is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that users from each company can view only the data of their respective company.

Why C is Wrong: a function is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 77

You have an Azure subscription.

You need to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. The solution must meet the following requirements:

* Minimize how long it takes to fail over.

* Maintain existing connections to the primary replica during a failover.

What should you do?

A. Connect each virtual machine to a single subnet on a single virtual network.

B. Connect each virtual machine to a single subnet on a virtual network. Deploy a standard Azure load balancer.

C. Connect each virtual machine to a different subnet on a single virtual network.

D. Connect each virtual machine to a different subnet on a virtual network. Deploy a basic Azure load balancer.

Answer: C (LEAVE A REPLY)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: Connect each virtual machine to a different subnet on a single virtual network. is correct because it is the feature whose normal purpose matches the stated requirement. Connect each virtual machine to a different subnet on a single virtual network. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You 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.

Why A is Wrong: Connect each virtual machine to a single subnet on a single virtual network. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: Connect each virtual machine to a single subnet on a virtual network. Deploy a standard Azure load balancer. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why D is Wrong: Connect each virtual machine to a different subnet on a virtual network. Deploy a basic Azure load balancer. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 78

Task 6

You need to ensure that you can connect to db1 by using a private IP address on a virtual network named VNET1. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

See the explanation part for the complete Solution.

Explanation:

Create an Azure Private Endpoint for the Azure SQL logical server that hosts db1, place the private endpoint in VNET1, and integrate it with the private DNS zone:

privatelink.database.windows.net

This is the correct solution because Azure SQL Database is a PaaS service. You do not assign a private IP directly to db1. Instead, Azure creates a private endpoint network interface in the virtual network. That private endpoint receives a private IP address from a subnet in VNET1, and clients in VNET1 use that private IP path to reach the SQL server. Microsoft defines a private endpoint as a network interface that uses a private IP address from your virtual network to connect privately to a Private Link resource such as Azure SQL Database.

Azure Portal Method - Recommended for Simulation

Step 1: Identify the SQL logical server that hosts db1

Sign in to the Azure portal.

Search for SQL databases.

Open db1.

On the database overview page, identify the Server name.

The private endpoint is created for the Azure SQL logical server, not for the database object alone. For Azure SQL Database, the Private Link resource type is:

Microsoft.Sql/servers

and the target subresource is:

sqlServer

Microsoft lists Azure SQL Database private endpoint DNS configuration under Microsoft.Sql/servers with subresource sqlServer.

Step 2: Open the SQL server networking page

Open the Azure SQL logical server that hosts db1.

In the left menu, go to:

Security > Networking

Select the Private access tab.

Select Create a private endpoint.

Microsoft's Azure SQL private endpoint workflow is performed from the SQL server resource under Networking > Private access, where you can create or manage private endpoint connections.

Step 3: Configure the private endpoint basics

On the Create private endpoint page:

Setting

Value

Subscription

Use the lab subscription

Resource group

Use the lab resource group

Name

pe-db1-sql

Region

Same region as VNET1, if possible

The name is not exam-critical. The critical part is that the endpoint is associated with VNET1 and the SQL server that hosts db1.

Step 4: Configure the target resource

On the Resource tab, configure:

Setting

Value

Connection method

Connect to an Azure resource in my directory

Resource type

Microsoft.Sql/servers

Resource

SQL logical server that hosts db1

Target sub-resource

sqlServer

Do not choose storage, VM, managed instance, or any unrelated resource type. This is Azure SQL Database, so the target subresource must be sqlServer.

Step 5: Configure VNET1 and subnet

On the Virtual Network tab:

Setting

Value

Virtual network

VNET1

Subnet

Select an available subnet in VNET1

Private IP configuration

Dynamic is fine unless the lab requires static

Azure will create a network interface for the private endpoint and assign it a private IP address from the selected subnet. Microsoft notes that the network interface page for the private endpoint shows the private IP address assigned to the private endpoint connection.

Step 6: Configure private DNS integration

On the DNS tab:

Enable private DNS zone integration.

Use or create the private DNS zone:

privatelink.database.windows.net

Link the private DNS zone to:

VNET1

This is not optional in a clean exam solution. Without DNS integration, clients may still resolve the SQL server name to the public endpoint instead of the private endpoint. Microsoft states that DNS is critical because it resolves the private endpoint IP address, and for Azure SQL Database the recommended private DNS zone is privatelink.database.windows.net.

Step 7: Review and create

Select Review + create.

Confirm:

Resource: SQL logical server hosting db1

Target subresource: sqlServer

Virtual network: VNET1

Private DNS zone: privatelink.database.windows.net

Select Create.

After deployment, the SQL server will have a private endpoint connection associated with VNET1.

Step 8: Approve the private endpoint connection if required

In most same-directory deployments, approval may be automatic. If approval is pending:

Open the SQL logical server.

Go to:

Networking > Private access

Select the pending private endpoint connection.

Select Approve.

Microsoft documents that SQL administrators can approve or reject private endpoint connections from the SQL server private access page.

Step 9: Optional but recommended - Disable public network access

The task only says you need to connect by private IP from VNET1. It does not explicitly say to block public access. But if the exam expects private-only access, then disable public access after the private endpoint works.

On the SQL logical server:

Go to:

Security > Networking > Public access

Set Public network access to:

Disabled

or select:

Deny public network access

Save.

Be careful: Microsoft states that adding a private endpoint does not automatically block public routing to the logical server. Public access must be denied separately if you want private-only access.

How to Connect from SSMS

You should connect from a machine that is inside VNET1, such as an Azure VM joined to VNET1.

Step 1: Test DNS from a VM in VNET1

From a VM in VNET1, run:

nslookup < sql-server-name > .database.windows.net

Expected result: the name should resolve through the private endpoint path and return a private IP address from VNET1's address space.

Microsoft explains that connection URLs do not change; DNS resolution is overridden so the existing service FQDN resolves to the private endpoint private IP address.

Step 2: Connect with SSMS

In SSMS, connect using the normal Azure SQL server name:

< sql-server-name > .database.windows.net

Then select database:

db1

Use normal SQL authentication or Microsoft Entra authentication.

Do not type the raw private IP address into SSMS unless the lab specifically forces it. For Azure SQL, the correct operational pattern is to connect to the SQL server FQDN and allow private DNS to resolve that FQDN to the private endpoint IP. Direct IP connection can cause TLS/certificate name problems because the server certificate matches the DNS name, not the private IP.

Verification

The task is complete when all of these are true:

Private endpoint exists for the SQL logical server hosting db1.

Target subresource is sqlServer.

The private endpoint is deployed into VNET1.

A private IP address is assigned to the private endpoint NIC.

Private DNS zone privatelink.database.windows.net exists.

The private DNS zone is linked to VNET1.

The SQL server FQDN resolves to the private endpoint private IP from inside VNET1.

SSMS can connect to db1 from a VM or client connected to VNET1.

Final Exam-Lab Action

Use the Azure portal and configure:

SQL server hosting db1

> Networking

> Private access

> Create private endpoint

Resource type: Microsoft.Sql/servers

Target subresource: sqlServer

Virtual network: VNET1

Private DNS zone: privatelink.database.windows.net

Then connect from a VM or client in VNET1 using:

< sql-server-name > .database.windows.net

That is the correct way to ensure db1 is reachable through a private IP address on VNET1.

NEW QUESTION: 79

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

Name	Type	Description
SQLSVR1	Logical SQL server	Hosts one Azure SQL database
SQLDB11	Azure SQL database	Is hosted by SQLSVR1
SQLSVR2	Logical SQL server	Hosts two Azure SQL databases
SQLDB21	Azure SQL database	Is hosted by SQLSVR2
SQLDB22	Azure SQL database	Is hosted by SQLSVR2

You plan louse SQLDB11 as an elastic job database to run jobs on SQIDb11 andSOtDB22. What is the minimum number of database scoped credentials required tor the elastic jobs?

- A. 1
- B. 2
- C. 3
- D. 4

Answer: (SHOW ANSWER)

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why B is Correct: 2 matches the expected DP-300 administration action. 2 can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: What is the minimum number of database scoped credentials required tor the elastic jobs?

Why A is Wrong: 1 can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why C is Wrong: 3 can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why D is Wrong: 4 can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

NEW QUESTION: 80

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines. The virtual machine hosts a database named DB1. You need to monitor DB1 by using Extended Events. The solution must meet the following requirements:

- * Capture raw event data and store the data in Azure Storage.
- * Minimize the performance impact of capturing extended events.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.


NOTE: Each correct selection is worth one point.

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
    ACTION (sqlserver.sql_text)
    WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
```

ADD TARGET

package0.



A dropdown menu showing four options: 'event_file', 'event_file', 'event_stream', and 'ring_buffer'. The first 'event_file' option is highlighted with a dark background.

)

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

)

WITH

```
(MAX_MEMORY = 10 MB,
```

```
EVENT_RETENTION_MODE=
```



A dropdown menu showing three options: 'ALLOW_MULTIPLE_EVENT_LOSS', 'ALLOW_SINGLE_EVENT_LOSS', and 'NO_EVENT_LOSS'. The 'ALLOW_MULTIPLE_EVENT_LOSS' option is highlighted with a dark background.

```
MAX_DISPATCH_LATENCY = 3 SEC
```

Answer:

Answer Area

```
CREATE EVENT SESSION session1 ON DATABASE
ADD EVENT sqlserver.sql_statement_starting
(
    ACTION (sqlserver.sql_text)
    WHERE statement LIKE 'UPDATE gmTabEmployee%'
)
ADD TARGET
package0. event_file
(
    SET filename = 'https://gmstorageaccountxevent.blob.core.windows.net/gmcontainerxevent/anyfilenamexel242b.xel'
)
WITH
    (MAX_MEMORY = 10 MB,
    EVENT_RETENTION_MODE= ALLOW_MULTIPLE_EVENT_LOSS
    MAX_DISPATCH_LATENCY = 3 SEC)
Microsoft
```

The image shows a SQL Server Enterprise Manager interface for creating an event session. The 'Event File' dropdown menu is open, showing options: event_file, event_stream, and ring_buffer. The 'Event Retention Mode' dropdown menu is also open, showing options: ALLOW_MULTIPLE_EVENT_LOSS, ALLOW_SINGLE_EVENT_LOSS, and NO_EVENT_LOSS. The 'Event File' dropdown is currently set to 'event_file' and the 'Event Retention Mode' dropdown is currently set to 'ALLOW_MULTIPLE_EVENT_LOSS'.

Explanation:

Answer Area

Microsoft

CREATE EVENT SESSION session1 ON DATABASE

ADD EVENT sqlserver.sql_statement_starting

(

ACTION (sqlserver.sql_text)

WHERE statement LIKE 'UPDATE gmTabEmployee%'

)

ADD TARGET

package0.

event_file

(

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

)

WITH

(MAX_MEMORY = 10 MB,

EVENT_RETENTION_MODE=

ALLOW_MULTIPLE_EVENT_LOSS

MAX_DISPATCH_LATENCY = 3 SECONDS)

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to monitor DB1 by using Extended Events.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

NEW QUESTION: 81

You have two Azure SQL databases named SQLDB1 and SQLDB2 in the Hyperscale service tier. SQLDB1 was set to the Hyperscale service tier when the database was created. SQLDB2 was changed from the Business Critical service tier to the Hyperscale service tier during the last 14 days.

You need to change SQLDB1 and SQLDB2 to the Business Critical service tier. The solution must minimize the administrative effort.

What should you do first? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

On SQLDB1:

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

On SQLDB2:

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

Answer:

NEW QUESTION

On SQLDB1:

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

On SQLDB2:

- Export the data.
- Set Service tier to Business Critical.
- Set Service tier to General Purpose.

Microsoft

Explanation:

Answer Area

On SQLDB1: Export the data.

On SQLDB2: Set Service tier to General Purpose.

Microsoft

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to change SQLDB1 and SQLDB2 to the Business Critical service tier.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

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.



To reduce how long it takes to complete the query you must [answer choice].

- scale the resource
- use an elastic pool
- perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

- a temporary
- an In-Memory OTLP durable
- an In-Memory OTLP non durable

Answer:

To reduce how long it takes to complete the query you must [answer choice].

- scale the resource
- use an elastic pool
- perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

- a temporary
- an In-Memory OTLP durable
- an In-Memory OTLP non durable

Explanation:

To reduce how long it takes to complete the query you must [answer choice].

scale the resource
use an elastic pool
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

a temporary
an In-Memory OTLP durable
an In-Memory OTLP non durable

Microsoft

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure SQL database named that contains a table named Table1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 83

You have two on-premises Microsoft SQL Server instances named SQL1 and SQL2.

You have an Azure subscription

You need to sync a subset of tables between the databases hosted on SQL1 and SQL2 by using SQL Data Sync.

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

Actions

Microsoft

Answer Area

freecram.net

Deploy an Azure SQL database.

Install and configure the Client Sync Agent app on SQL1 and SQL2.

Create a sync group.

Configure the sync group.

Sync the metadata database configuration.

Deploy an Azure SQL managed instance.

Answer:

Actions

Microsoft

Answer Area

freecram.net

Deploy an Azure SQL database.

Install and configure the Client Sync Agent app on SQL1 and SQL2.

Create a sync group.

Configure the sync group.

Sync the metadata database configuration.

Deploy an Azure SQL managed instance.

Install and configure the Client Sync Agent app on SQL1 and SQL2.

Create a sync group.

Configure the sync group.

Sync the metadata database configuration.

Deploy an Azure SQL managed instance.

Explanation:

Actions

⋮

Deploy an Azure SQL database.

Answer Area

1

⋮

Install and configure the Client Sync Agent app on SQL1 and SQL2.

2

⋮

Create a sync group.

3

⋮

Configure the sync group.

4

⋮

Sync the metadata database configuration.

5

⋮

Deploy an Azure SQL managed instance.

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure subscription You need to sync a subset of tables between the databases hosted on SQL1 and SQL2 by using SQL Data Sync.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 84

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 area.

NOTE: Each correct selection is worth one point.

To collect monitoring data, use:

A virtual machine

An Azure function

The Azure Monitor agent

To store monitoring data, create:

A Log Analytics workspace

An Azure SQL database

An Azure Storage account

Answer:

To collect monitoring data, use:

To store monitoring data, create:



Explanation:

Box 1 = Azure Monitor Agent

Box 2 = An Azure SQL database

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to monitor DB 1 by using SQL Insights.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

NEW QUESTION: 85

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

Name	Type	Configuration
DB1	Azure SQL Database	Hyperscale service tier No secondary replicas
App1	Azure Web Apps	App1 has read-only access to DB1. There are multiple instances of App1.

You need to create a read-only replica of DB1 and configure the App1 instances to use the replica.

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

NOTE: Each correct selection is worth one point.

Answer Area

To add read-only replicas of DB1:

Create a replica on the same logical server.
Create a new logical server and configure geo-replication.
Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

Add an ApplicationIntent entry to the connection string.
Add a MultiSubnetFailover entry to the App1 connection string.
Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Answer:

Answer Area

To add read-only replicas of DB1:

Create a replica on the same logical server. |
Create a new logical server and configure geo-replication.
Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

Add an ApplicationIntent entry to the connection string. |
Add a MultiSubnetFailover entry to the App1 connection string.
Create a dedicated endpoint and configure the App1 connection string to point to the endpoint.

Explanation:

To add read-only replicas of DB1:

▼

Create a replica on the same logical server.

Create a new logical server and configure geo-replication.

Create a new logical server and configure an auto-failover group.

To configure App1 instances to access the read-only replica:

Add an ApplicationIntent entry to the connection string.

Add a MultiSubnetFailover entry to the App1 connection string.

Create a dedicated endpoint and configure the App1 connection string to point to the

Reference:
<https://sqlserverguides.com/read-only-replica-azure-sql/>
Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.
Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a read-only replica of DB1 and configure the App1 instances to use the replica.
Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 86

You need to deploy an Azure SQL Database elastic pool by using a Bicep template.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
resource symbolicname 'Microsoft.Sql/elasticPools@2021-11-01-preview' = {  
  name: 'string'  
  location: 'string'  
  tags: {  
  }  
  sku: {  
    ...  
  }  
  parent: resourceSymbolicName  
  properties: {  
    highAvailabilityReplicaCount: int  
    licenseType: 'string'  
    maintenanceConfigurationId: 'string'  
    maxSizeBytes: int  
    capacity : int  
    {  
      Capacity  
      maxCapacity  
      MinCapacity  
      maintenanceConfigurationId  
    }  
    zoneRedundant: bool  
  }  
}
```




Answer:

Answer Area



```
resource symbolicname 'Microsoft.Sql/elasticPools@2021-11-01-preview' = {  
  name: 'string'  
  location: 'string'  
  tags: {  
  }  
  sku: {  
    ...  
  }  
  parent: resourceSymbolicName  
  properties: {  
    highAvailabilityReplicaCount: int  
    licenseType: 'string'  
    maintenanceConfigurationId: 'string'  
    maxSizeBytes: int  
    Capacity: {  
      maxCapacity: int  
      minCapacity: int  
    }  
    zoneRedundant: bool  
  }  
}
```

Explanation:



Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to deploy an Azure SQL Database elastic pool by using a Bicep template.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 87

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: (SHOW ANSWER)

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

[tabs=azure-portal](#)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why B and C is Correct: Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. In this scenario, the important constraint is: You need to ensure that users can connect to DB1 in the event of an Azure regional outage. From the properties of Server1 add a failover group. satisfies that constraint without adding an unrelated service or manual process. Create a new Azure SQL Database server named Server2.

matches the expected DP-300 administration action. Create a new Azure SQL Database server named Server2. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to ensure that users can connect to DB1 in the event of an Azure regional outage.

Why A is Wrong: Active geo-replication creates readable secondary databases, but failover groups add the managed listener and automatic failover abstraction needed for application continuity at scale. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: From the properties of Server1 configure retention for DB1 is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why E is Wrong: Create a new Azure SQL Database instance named DB2. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 88


You have an instance of SQL Server on Azure Virtual Machines named VM1.

You need to use an Azure Automation runbook to initiate a SQL Server database backup on VM1.

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



Explanation:



Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to use an Azure Automation runbook to initiate a SQL Server database backup on VM1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 89

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: (SHOW ANSWER)

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: 90

You have an Azure subscription.

You need to deploy a logical SQL server by using an Azure Resource Manager (ARM) template. The solution must ensure that the server will allow inbound connectivity from any Azure resource.

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

Microsoft

```
"resources": [
  {
    "type": "Microsoft.Sql/servers",
    ...
  },
  {
    "type": "Microsoft.Sql/servers/firewallRules",
    "apiVersion": "2021-02-01-preview",
    "name": "[concat(parameters('servers_a400n10102_name'), '/AllowAllWindowsAzureIps')]",
    "dependsOn": [
      "[resourceId('Microsoft.Sql/servers', parameters('servers_a400n10102_name'))]"
    ],
    "properties": {
      "startIpAddress": "0.0.0.0",
      "endIpAddress": "0.0.0.0"
    }
  }
]
```

start

0.0.0.0

255.255.255.255

end

0.0.0.0

255.255.255.255

Answer:
Answer Area

Microsoft

```
"resources": [
  {
    "type": "Microsoft.Sql/servers",
    ...
  },
  {
    "type": "Microsoft.Sql/servers/firewallRules",
    "apiVersion": "2021-02-01-preview",
    "name": "[concat(parameters('servers_a400n10102_name'), '/AllowAllWindowsAzureIps')]",
    "dependsOn": [
      "[resourceId('Microsoft.Sql/servers', parameters('servers_a400n10102_name'))]"
    ],
    "properties": {
      "startIpAddress": "0.0.0.0",
      "endIpAddress": "0.0.0.0"
    }
  }
]
```

start

0.0.0.0

255.255.255.255

end

0.0.0.0

255.255.255.255

Explanation:

QUESTION 91

```
"resources": [
  {
    "type": "Microsoft.Sql/servers",
    ...
  },
  {
    "type": "Microsoft.Sql/servers/firewallRules",
    "apiVersion": "2021-02-01-preview",
    "name": "[concat(parameters('servers_e400m10102_name'), '/AllowAllWindowsAzureIps')]",
    "dependsOn": [
      "[resourceId('Microsoft.Sql/servers', parameters('servers_e400m10102_name'))]"
    ],
    "properties": {
      "startIpAddress": "0.0.0.0",
      "endIpAddress": "0.0.0.0"
    }
  }
]
```

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to deploy a logical SQL server by using an Azure Resource Manager (ARM) template.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 91

Task 4

You need to encrypt the LastName column of the SalesLT.Customer table in db1 by using Always Encrypted.

You must use the Windows Certificate Store. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

See the explanation part for the complete Solution.

Explanation:

Use SQL Server Management Studio Always Encrypted Wizard and encrypt:

Database: db1

Schema: SalesLT

Table: Customer

Column: LastName

Encryption feature: Always Encrypted

Column master key store: Windows Certificate Store

Do not use Transparent Data Encryption. TDE encrypts the database at rest, but this task specifically requires Always Encrypted, which protects selected columns and keeps the encryption keys outside the database engine. Microsoft states that Always Encrypted uses column encryption keys to encrypt column data and column master keys to protect those column encryption keys. Column master keys are stored outside the database system, such as in the Windows certificate store.

Method 1 - SSMS Always Encrypted Wizard

This is the correct method for the simulation.

Step 1: Open SSMS and connect to Azure SQL Database

Open SQL Server Management Studio.

Connect to the Azure SQL logical server that hosts db1.

Use a SQL admin account or Microsoft Entra admin account.

In Options > Connection Properties, select database:

db1

Connect.

If SSMS cannot connect, go to the Azure portal and add your client IP address under the SQL server firewall /networking settings.

Step 2: Open the table column

In Object Explorer:

Expand Databases.

Expand db1.

Expand Tables.

Expand:

SalesLT.Customer

Expand Columns.

Locate:

LastName

Microsoft confirms that the Always Encrypted Wizard can be launched at the database, table, or individual column level. For one column, the cleanest path is to launch it directly from the column.

Step 3: Launch the Always Encrypted Wizard

Right-click the LastName column, then select:

Encrypt Column

or, depending on the SSMS version:

Always Encrypted Wizard

Alternative path:

Right-click db1 > Tasks > Always Encrypted Wizard

Then manually select:

SalesLT.Customer.LastName

Step 4: Select the LastName column for encryption

On the Column Selection page:

Find:

SalesLT.Customer.LastName

Select the checkbox for LastName.

Set the encryption type.

Use:

Randomized

unless the lab specifically requires searching or equality filtering on LastName.

Reason: Randomized encryption is stronger because identical plaintext values produce different ciphertext values. Deterministic encryption allows equality lookups, joins, grouping, and indexing, but leaks more pattern information because identical plaintext values produce identical encrypted values. Microsoft describes deterministic encryption as query-friendly but more pattern-revealing, while randomized encryption is more secure but does not support normal searching/grouping/joining without secure enclaves.

For this task, the requirement is only to encrypt the LastName column, so Randomized is the safer default.

Step 5: Choose or create a Column Encryption Key

For Encryption Key, select a new key such as:

CEK_Auto1

or create a new column encryption key if one does not already exist.

This is the key that encrypts the data in the LastName column. Microsoft states that a column encryption key encrypts the data in encrypted columns, and the column master key encrypts/protects the column encryption key.

Select Next.

Step 6: Configure the Column Master Key in Windows Certificate Store

On the key configuration page, create a new Column Master Key.

Use settings like:

Setting

Value

Column master key name

CMK_WindowsCert or default generated name

Key store

Windows Certificate Store

Certificate location

Current User

Certificate

Generate new certificate

Column encryption key

CEK_Auto1 or default generated CEK

In many SSMS versions, the wizard creates both:

CMK_Auto1

CEK_Auto1

That is acceptable as long as the CMK key store is Windows Certificate Store.

Microsoft states that SQL Server Management Studio supports column master keys stored in the Windows Certificate Store, and that a column master key can be a certificate stored in Windows Certificate Store.

Step 7: Run the wizard

On the final wizard pages:

Review the configuration.

Choose:

Proceed to finish now

or:

Run immediately

Select Finish.

SSMS will generate the column master key metadata, column encryption key metadata, and perform the data encryption operation. Microsoft explains that the wizard can encrypt selected columns and can generate a new column master key and column encryption key when needed.

During encryption, SSMS may temporarily create a new table, copy data, encrypt the selected column, and swap the table back, depending on whether secure enclaves are used. Microsoft notes that the wizard can move data out of the database and perform cryptographic operations inside the SSMS process when secure enclave in-place encryption is not used.

Verification Steps

Step 1: Confirm Always Encrypted metadata exists

Run this in db1:

SELECT

```
name,  
key_store_provider_name,  
key_path  
FROM sys.column_master_keys;
```

You should see a column master key that uses the Windows certificate store provider.

Then run:

```
SELECT  
name  
FROM sys.column_encryption_keys;
```

You should see the column encryption key created by the wizard.

Step 2: Confirm LastName is encrypted

Run:

```
SELECT  
t.name AS table_name,  
c.name AS column_name,  
c.encryption_type_desc,  
cek.name AS column_encryption_key  
FROM sys.columns AS c  
JOIN sys.tables AS t  
ON c.object_id = t.object_id  
LEFT JOIN sys.column_encryption_keys AS cek  
ON c.column_encryption_key_id = cek.column_encryption_key_id  
WHERE t.name = ' Customer '  
AND SCHEMA_NAME(t.schema_id) = ' SalesLT '  
AND c.name = ' LastName ' ;
```

Expected result:

table_name: Customer

column_name: LastName

encryption_type_desc: RANDOMIZED

column_encryption_key: CEK_Auto1 or similar

If you selected deterministic encryption, the expected value will be:

DETERMINISTIC

The key requirement is that encryption_type_desc is no longer NULL.

Step 3: Test viewing the encrypted column

Open a new SSMS query connection without Always Encrypted enabled and run:

```
SELECT TOP 10 LastName  
FROM SalesLT.Customer;
```

You should not see normal plaintext values.

Then reconnect with Always Encrypted enabled:

In SSMS, select Connect > Database Engine.

Select Options.

Go to Additional Connection Parameters.

Add:

Column Encryption Setting=Enabled

Connect again.

Run:

SELECT TOP 10 LastName

FROM SalesLT.Customer;

A client that has access to the Windows certificate/private key should be able to decrypt the values. Microsoft explains that the database stores encrypted data and key metadata, while client-side components with access to the column master key perform encryption and decryption.

Important Exam Notes

Do not choose Azure Key Vault

The task explicitly says:

You must use the Windows Certificate Store.

So the column master key should not be stored in Azure Key Vault.

Wrong:

Azure Key Vault

Correct:

Windows Certificate Store

Do not use TDE

TDE is not column-level Always Encrypted. It encrypts database files/logs at rest, but users and administrators querying the database still see plaintext if they have SQL permissions.

Correct technology:

Always Encrypted

Correct key store:

Windows Certificate Store

Correct target column:

SalesLT.Customer.LastName

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (**503** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 92

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:

The Server Audit Specification object belongs to an audit.

A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.

Reference:

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

NEW QUESTION: 93

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. SQLAgentUserRole

C. SQLAgentReaderRole

D. SQLAgentOperatorRole

Answer: A ([LEAVE A REPLY](#))

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why A is Correct: sysadmin is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. This fits the case because the requested outcome is: You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs. The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: SQLAgentUserRole is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why C is Wrong: SQLAgentReaderRole is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: SQLAgentOperatorRole is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs.

NEW QUESTION: 94

You have five on-premises servers that have Microsoft SQL Server 2022 Enterprise installed. Each server contains multiple databases that store data for web apps and are backed up by using a third-party backup solution.

You plan to migrate the databases to Azure.

You need to recommend a solution to host the databases. The solution must meet the following requirements:

* Compute and storage resources must be shared across the databases.

* Costs must be minimized.

What should you include in the recommendation?

A. SQL Server on Azure Virtual Machines

B. in Azure SQL Database elastic pool

C. Azure SQL Database

D. Azure SQL Managed Instance

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. In this scenario, the important constraint is: You have five on-premises servers that have Microsoft SQL Server 2022 Enterprise installed Each server contains multiple database that store data for web apps and are backed up by using a third-party backup solution. in Azure SQL Database elastic pool satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: SQL Server on Azure Virtual Machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Azure SQL Database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: Azure SQL Managed Instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

NEW QUESTION: 95

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: ([SHOW ANSWER](#))

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-database-portable?view=sql-server-ver15> Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why D is Correct: a contained database user from a Windows login for App1 on DB2 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to ensure that App1 can access DB2. a contained database user from a Windows login for App1 on DB2 satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: a contained database user for App1 on DB2 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: a login for App1 on Server1 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that App1 can access DB2.

Why C is Wrong: a contained database user from an external provider for App1 on DB2 is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 96

You have a data warehouse in Azure Synapse Analytics.

You need to ensure that the data in the data warehouse is encrypted at rest.

What should you enable?

- A.** Transparent Data Encryption (TDE)
- B.** Advanced Data Security for this database
- C.** Always Encrypted for all columns
- D.** Secure transfer required

Answer: ([SHOW ANSWER](#))

Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-tde-overview> Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why A is Correct: Transparent Data Encryption (TDE) is correct because it is the feature whose normal purpose matches the stated requirement. Transparent Data Encryption protects database files, backups, and logs at rest; it does not encrypt data in use or enforce user-level data visibility. The scenario wording points to that specific behavior: You need to ensure that the data in the data warehouse is encrypted at rest.

Why B is Wrong: Advanced Data Security for this database is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why C is Wrong: Always Encrypted protects sensitive columns from exposure to the database engine by encrypting and decrypting data at the client side with keys protected outside the database. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that the data in the data warehouse is encrypted at rest.

Why D is Wrong: Secure transfer required is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 97

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.

You need to ensure that each customer can only connect to and access their respective database.

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

NOTE: Each correct selection is worth one point.

- A.** Implement row-level security (RLS).
- B.** Create users in each database.
- C.** Configure the database firewall.
- D.** Configure the server firewall.
- E.** Create logins in the master database.
- F.** Implement Always Encrypted.

Answer: ([SHOW ANSWER](#))

Manage database access by adding users to the database, or allowing user access with secure connection strings.

Database-level firewall rules only apply to individual databases.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/secure-database-tutorial> Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B and E is Correct: Create users in each database. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. This fits the case because the requested outcome is: You need to ensure that each customer can only connect to and access their respective database. The selected feature addresses that outcome directly rather than relying on a workaround. Create logins in the master database. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to ensure that each customer can only connect to and access their respective database. Create logins in the master database. satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Row-Level Security filters rows returned by queries based on a predicate function, allowing different users or tenants to see different subsets of the same tables. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why C is Wrong: Database-level firewall rules restrict access to a specific Azure SQL database and are more granular than server-level firewall rules. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why D is Wrong: Server-level firewall rules apply to all databases on the logical server and are broader than database-level firewall rules. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why F is Wrong: Always Encrypted protects sensitive columns from exposure to the database engine by encrypting and decrypting data at the client side with keys protected outside the database. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 98

You have an Azure Synapse Analytics workspace named WS1 that contains an Apache Spark pool named Pool1.

You plan to create a database named DB1 in Pool1.

You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool.

Which format should you use for the tables in DB1?

A. JSON

B. CSV

C. Parquet

D. ORC

Answer: ([SHOW ANSWER](#))

Serverless SQL pool can automatically synchronize metadata from Apache Spark. A serverless SQL pool database will be created for each database existing in serverless Apache Spark pools.

For each Spark external table based on Parquet and located in Azure Storage, an external table is created in a serverless SQL pool database. As such, you can shut down your Spark pools and still query Spark external tables from serverless SQL pool.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-storage-files-spark-tables> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why C is Correct: Parquet is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to ensure that when tables are created in DB1, the tables are available automatically as external tables to the built-in serverless SQL pool. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: JSON is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: CSV is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: ORC is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 99

You need to identify the cause of the performance issues on SalesSQLDb1.

Which two dynamic management views should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. sys.dm_pdw_nodes_tran_locks

B. sys.dm_exec_compute_node_errors

C. sys.dm_exec_requests

D. sys.dm_cdc_errors

E. sys.dm_pdw_nodes_os_wait_stats

F. sys.dm_tran_locks

Answer: ([SHOW ANSWER](#))

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information.

```
SELECT
t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
```

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

sys.dm_os_wait_stats: for SQL Server

sys.dm_db_wait_stats: for Azure SQL Database

sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql>

Topic 1, Litware

Existing Environment

Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering.

The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

Database Environment

The sales department has the following database workload:

An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.

A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1.

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1. Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.

ResearchDB1 will contain Personally Identifiable Information (PII) data.

Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

The use of an on-premises infrastructure must be minimized.

Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

Encrypt the PII data at rest, in transit, and in use.

Use the principle of least privilege whenever possible.

Authenticate database users by using Active Directory credentials.

Protect Azure SQL Database instances by using database-level firewall rules.

Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

Business Requirements

Litware identifies the following business requirements:

Meet an SLA of 99.99% availability for all Azure deployments.

Minimize downtime during the migration of the SERVER1 databases.

Use the Azure Hybrid Use Benefits when migrating workloads to Azure.

Once all requirements are met, minimize costs whenever possible.

NEW QUESTION: 100

You have an Azure subscription that contains the following resources:

- * 10 Azure SQL databases

- * Five Azure SQL managed instances

- * Five instances of SQL Server on Azure Virtual Machines

You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort. What should you include in the solution?

A. Log Analytics

B. Azure SQL Analytics

C. Query Performance Insight

D. SQL Insights

Answer: (SHOW ANSWER)

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why B is Correct: Azure SQL Analytics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You have an Azure subscription that contains the following resources: * 10 Azure SQL databases * Five Azure SQL managed instances * Five instances of SQL Server on Azure Virtual Machines You need to implement a centralized monitoring solution for all the Azure The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Log Analytics stores logs and metrics in a queryable workspace and supports workbooks, dashboards, alerting, and cross-resource analysis. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Query Performance Insight surfaces resource-consuming queries and database workload patterns in Azure SQL for performance investigation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: SQL Insights is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 101

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: (SHOW ANSWER)

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: 102

You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.

You need to recommend a deployment solution that meets the following requirements:

- * Provides a Service Level Agreement (SLA) of at least 99.95%
- * Replicates databases in the same group synchronously
- * Minimizes the latency of database writes

What should you recommend?

- A.** Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximitygroup.
- B.** Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group.
- C.** Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group.
- D.** Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group.

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why A is Correct: Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximitygroup. is correct because it is the feature whose normal purpose matches the stated requirement. Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximitygroup. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to recommend a deployment solution that meets the following requirements: * Provides a Service Level Agreement (SLA) of at least 99.95% * Replicates databases in the same group synchronously * Minimizes the latency of database writes What should you r Why B is Wrong: Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 103

You have an Azure subscription.

You need to deploy an Azure SQL database. The solution must meet the following requirements:

- * Dynamically scale CPU resources.
- * Ensure that the database can be paused to reduce costs.

What should you use?

- A.** the Business Critical service tier
- B.** the serverless compute tier
- C.** an elastic pool
- D.** the General Purpose service tier

Answer: (SHOW ANSWER)

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why B is Correct: the serverless compute tier is correct because it is the feature whose normal purpose matches the stated requirement. The serverless tier automatically scales compute and can pause during inactive periods, making it suitable for intermittent workloads with bursts and idle windows. The scenario wording points to that specific behavior: You need to deploy an Azure SQL database.

Why A is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.


Why C is Wrong: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why D is Wrong: the General Purpose service tier is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

NEW QUESTION: 104

You have an Azure SQL database named DB1 that has default configurations.

You review the Data Discovery & Classification blade for D61 as shown in the following exhibit.



You need to identify the process for creating a custom sensitivity label and assigning it to the CustomerID column. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point

- A. Select Export
- B. Filter by the Contact Info Information type-
- C. Filter by the CustomerID column name.
- D. Select the Classification page.

E. Select Configure

Answer: (SHOW ANSWER)

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why D and E is Correct: Select the Classification page. is correct because it is the feature whose normal purpose matches the stated requirement. Select the Classification page. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario wording points to that specific behavior: You need to identify the process for creating a custom sensitivity label and assigning it to the CustomerID column. Select Configure is correct because it is the feature whose normal purpose matches the stated requirement. Select Configure is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario wording points to that specific behavior: You need to identify the process for creating a custom sensitivity label and assigning it to the CustomerID column.

Why A is Wrong: Select Export is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why B is Wrong: Filter by the Contact Info Information type- is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: Filter by the CustomerID column name. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

NEW QUESTION: 105

You have a burstable Azure virtual machine named VM1 that hosts an instance of Microsoft SQL Server.

You need to attach an Azure ultra disk to VM1. The solution must minimize downtime on VM1.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Attach the ultra disk.

Stop and deallocate VM1.

Set Enable Ultra disk compatibility to **Yes**.

Resize VM1.

Start VM1.

Answer Area

Microsoft

freecram.net

>

<

Answer:

Actions

Attach the ultra disk.

Stop and deallocate VM1.

Set Enable Ultra disk compatibility to **Yes**.

Resize VM1.

Start VM1.

Answer Area

Stop and deallocate VM1.

Attach the ultra disk.

➤ Set Enable Ultra disk compatibility to **Yes**.

⬅ Resize VM1.

Start VM1.

Explanation:

Answer Area

1 Stop and deallocate VM1.

2 Attach the ultra disk.

3 Set Enable Ultra disk compatibility to **Yes**.

4 Resize VM1.

5 Start VM1.

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to attach an Azure ultra disk to VMI.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 106

You plan to deploy an Always On failover cluster instance (FCI) on Azure virtual machines.

You need to provision an Azure Storage account to host a cloud witness for the deployment.

How should you configure the storage account? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Type: Standard general-purpose v2
Premium block blobs
Premium file shares
Standard general-purpose v2

Replication: Locally-redundant storage (LRS)
Geo-redundant storage (GRS)
Locally-redundant storage (LRS)
Read-access geo-redundant storage (RA-GRS)

Microsoft

Answer:
Answer Area

Microsoft

Type: Standard general-purpose v2
Premium block blobs
Premium file shares
Standard general-purpose v2

Replication: Locally-redundant storage (LRS)
Geo-redundant storage (GRS)
Locally-redundant storage (LRS)
Read-access geo-redundant storage (RA-GRS)

Explanation:

Answer Area

Type: Standard general-purpose v2

Replication: Locally-redundant storage (LRS)

Microsoft

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to provision an Azure Storage account to host a cloud witness for the deployment.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (**503** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 107

You have SQL Server on an Azure virtual machine that contains a database named Db1.

You need to enable automatic tuning for Db1.

How should you complete the statements? To answer, select the appropriate answer in the answer area.

NOTE: Each correct selection is worth one point.

ALTER DATABASE [Db1]

▼

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)

SET AUTOMATIC_TUNING=AUTO

SET QUERY_STORE=OFF

SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)

SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO

ALTER DATABASE [Db1]

▼

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)

SET AUTOMATIC_TUNING=AUTO

SET QUERY_STORE=OFF

SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)

SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO

Answer:

ALTER DATABASE [Db1]

▼

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)

SET AUTOMATIC_TUNING=AUTO

SET QUERY_STORE=OFF

SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)

SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO

ALTER DATABASE [Db1]

▼

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF)

SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)

SET AUTOMATIC_TUNING=AUTO


SET QUERY_STORE=OFF

SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)

SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)

GO

Explanation:



Box 1: SET AUTOMATIC_TUNING = AUTO

To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

```
ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO
```

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

```
ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)
```

Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option.

Setting it to OFF will also override any setting that database inherited and disable the tuning option.

Reference:

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

NEW QUESTION: 108

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: (SHOW ANSWER)

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> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D is Correct: Azure SQL Analytics alerts based on diagnostics logs is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to configure alerts for each database based on the diagnostics telemetry of the database. Azure SQL Analytics alerts based on diagnostics logs satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Azure SQL Analytics alerts based on metrics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: SQL Health Check alerts based on diagnostics logs is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: SQL Health Check alerts based on metrics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 109

You have two resource groups named RG1 and RG2. RG1 contains an Azure SQL managed instance named MM that hosts a database named DB1. RG2 contains a SQL managed instance named MI2.

You run the following PowerShell script.

```
$database = 'DB1'
$sourceMI = 'MI1'
$sourceRG = 'RG1'
$targetMI = 'MI2'
$targetRG = 'RG2'

## Start database operation.
$databaseCopyParams = @(
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
)
Copy-AzSqlInstanceDatabase @databaseCopyParams

## Verify the operation status is succeeded.
}
Get-AzSqlInstanceDatabaseCopyOperation @databaseCopyOperationsParams

## Complete database operation.
$completeDatabaseCopyParams = @(
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
)
Complete-AzSqlInstanceDatabaseCopy @completeDatabaseCopyParams
```

Answer Area

Microsoft

Statements

	Yes	No
The copy of DB1 will be dropped from MI1.	☐	☐
A copy of DB1 on MI1 will be available during the operation.	☐	☐
A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.	☐	☐

Answer:

Answer Area

Microsoft

Statements

	Yes	No
The copy of DB1 will be dropped from MI1.	☐	☒
A copy of DB1 on MI1 will be available during the operation.	☒	☐
A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.	☐	☒

Explanation:

Answer Area

Microsoft

Statements

	Yes	No
The copy of DB1 will be dropped from MI1.	☐	☒
A copy of DB1 on MI1 will be available during the operation.	☒	☐
A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.	☐	☒

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have two resource groups named RG1 and RG2.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 110

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: 111

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: (SHOW ANSWER)

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> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why C is Correct: Assign a larger resource class to the automated data load queries. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait /blocking detail, or automatic remediation. This fits the case because the requested outcome is: You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Assign a smaller resource class to the automated data load queries. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait /blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Create sampled statistics to every column in each table of DW1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait /blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Hash distribute the large fact tables in DW1 before performing the automated data loads. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 112

You have an Azure SQL Database failover group named AzFG1.

You plan to run the following CLI command.

```
az sql failover-group update \
  --name AzFG1 \
  --grace-period 2 \
  --failover-policy Automatic
```

How will this affect AzFG1?

- A.** Up to two hours of data will be lost during an automatic failover.
- B.** Automatic failover will occur after a minimum of two hours.
- C.** A forced failover cannot be initiated during the grace period.
- D.** AzFG1 will attempt to fail back within two hours.

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why B is Correct: Automatic failover will occur after a minimum of two hours. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. In this scenario, the important constraint is: You have an Azure SQL Database failover group named AzFG1. Automatic failover will occur after a minimum of two hours. satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Up to two hours of data will be lost during an automatic failover. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: A forced failover cannot be initiated during the grace period. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: AzFG1 will attempt to fail back within two hours. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

NEW QUESTION: 113

You are working on two projects named Project1 and Project2 that will each migrate an on-premises app to Azure. Project1 must meet the following requirements:

- * Have access to the operating system of the database server that hosts the app.
- * Have 25 TB of storage.

Project2 must meet the following requirements:

- * Have change data capture (CDC) enabled
- * Support instance-level SQL auditing.
- * Support managed backups.
- * Have 10 TB of storage.
- * Offload reports.

You need to recommend an Azure SQL solution for the projects. The solution must minimize costs.

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

NOTE: Each correct selection is worth one point.

Answer Area

Project1:

Azure Kubernetes Service (AKS)

Azure SQL Database

Azure SQL Managed Instance in the Business Critical service tier

Azure SQL Managed Instance in the General Purpose service tier

SQL Server on Azure Virtual Machines

Project2:

Azure Kubernetes Service (AKS)

Azure SQL Database

Azure SQL Managed Instance in the Business Critical service tier

Azure SQL Managed Instance in the General Purpose service tier

SQL Server on Azure Virtual Machines

Answer:

Answer Area

Project1:

Azure Kubernetes Service (AKS)

Azure SQL Database

Azure SQL Managed Instance in the Business Critical service tier

Azure SQL Managed Instance in the General Purpose service tier

SQL Server on Azure Virtual Machines

Project2:

Azure Kubernetes Service (AKS)

Azure SQL Database

Azure SQL Managed Instance in the Business Critical service tier

Azure SQL Managed Instance in the General Purpose service tier

SQL Server on Azure Virtual Machines

Explanation:

Answer Area

Project1:

SQL Server on Azure Virtual Machines

Project2:

Azure SQL Managed Instance in the General Purpose service tier

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: Project1 must meet the following requirements: *

Have access to the operating system of the database server that hosts the app.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

NEW QUESTION: 114

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.

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

NOTE: Each correct selection is worth one point.

- 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
	7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	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: 115

You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.

You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.

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

Actions

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

Answer Area



Answer: Actions

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

Answer Area

- Create a resource pool.
- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.

Explanation:

- Create a resource pool.
- Create a workload group.
- Create a user-defined classifier function.

Modify Resource Governor

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15>

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-defined-function?view=sql-server-ver15> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 116

You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a virtual machine named Server1 that runs Windows Server. SQLVM1 and Server1 are joined to an Active Directory Domain Services (AD DS) domain. Server1 hosts a file share named Share1.

You need to ensure that a SQL Server Agent job step on SQLVM1 can access the files in Share1. The solution must use the principle of least privilege.

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

Actions

Create a login.

Create a database user.

Create a credential.

Create a proxy.

Assign the proxy to the job step.

Answer Area

Answer:

Actions

Create a login.

Create a database user.

Create a credential.

Create a proxy.

Assign the proxy to the job step.

Answer Area

Create a credential.

Create a proxy.

Assign the proxy to the job step.

Explanation:

Actions

Create a login.

Create a database user.

Answer Area

1 Create a credential.

2 Create a proxy.

3 Assign the proxy to the job step.

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that a SQL Server Agent job step on SQLVM1 can access the files in Share1.

Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

NEW QUESTION: 117

You have an Azure Stream Analytics job.

You need to ensure that the job has enough streaming units provisioned.

You configure monitoring of the SU % Utilization metric.

Which two additional metrics should you monitor? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Late Input Events
- B. Out of order Events
- C. Backlogged Input Events
- D. Watermark Delay
- E. Function Events

Answer: (SHOW ANSWER)

To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.

Note: 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: 118

Name	Location	Subnet
Branch1	Sydney	131.107.10.0/24
Branch2	Melbourne	131.107.11.0/24

You have an Azure subscription that contains an Azure SQL Database server named SQL1. SQL1 contains two databases named DB1 and DB2. You need to configure access to DB1 and DB2. The solution must meet the following requirements:

- * Ensure that DB1 can be accessed only by users in Branch1.
- * Ensure that DB2 can be accessed only by users in Branch2.

Solution: You connect to the master of SQL1 and run the following command.

execute sp_set_firewall_rule ' Allow dbl and db2 users " , ' 131.107.11.0 ' , " 131.107.11.255 ' Does this meet the goal?

- A. Yes
- B. No

Answer: (SHOW ANSWER)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: No is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario asks for: You need to configure access to DB1 and DB2. That makes No the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Yes is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data- platform design than the one the question is testing.

NEW QUESTION: 119

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A and E is Correct: Change Salary to the money data type. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to recommend changes to the data types to reduce storage and improve performance. Change Salary to the money data type. satisfies that constraint without adding an unrelated service or manual process. Change LastHireDate to the date data type.

is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to recommend changes to the data types to reduce storage and improve performance. Change LastHireDate to the date data type. satisfies that constraint without adding an unrelated service or manual process.

Why B is Wrong: Change PhoneNumber to the float data type. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Change LastHireDate to the datetime2(7) data type. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: Change PhoneNumber to the bigint data type. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 120

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 2:8:905856.

You need to improve system performance.

Solution: You shrink the transaction log file.

Does this meet the goal?

data-b

A. Yes

B. No

Answer: (SHOW ANSWER)

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why B is Correct: No is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: Each question in the series contains a unique solution that might meet the stated goals. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 121

You plan to create a table in an Azure Synapse Analytics dedicated SQL pool.

Data in the table will be retained for five years. Once a year, data that is older than five years will be deleted.

You need to ensure that the data is distributed evenly across partitions. The solutions must minimize the amount of time required to delete old data.

How should you complete the Transact-SQL statement? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

CustomerKey
HASH
ROUND_ROBIN
REPLICATE
OrderDateKey
SalesOrderNumber

Answer Area

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [ ] ([ProductKey])
, PARTITION ( [ ] RANGE RIGHT FOR VALUES
                (20170101, 20180101, 20190101, 20200101, 20210101)
            )
)
```

Answer:

Values

CustomerKey
HASH
ROUND_ROBIN
REPLICATE
OrderDateKey
SalesOrderNumber

Answer Area

```
CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = [ HASH ] ([ProductKey])
, PARTITION ( [ OrderDateKey ] RANGE RIGHT FOR VALUES
                (20170101, 20180101, 20190101, 20200101, 20210101)
            )
)
```

Explanation:

```

CREATE TABLE [dbo].[FactSales]
(
    [ProductKey] int NOT NULL
, [OrderDateKey] int NOT NULL
, [CustomerKey] int NOT NULL
, [SalesOrderNumber] nvarchar ( 20 ) NOT NULL
, [OrderQuantity] smallint NOT NULL
, [UnitPrice] money NOT NULL
)
WITH
(
    CLUSTERED COLUMNSTORE INDEX
, DISTRIBUTION = HASH ([ProductKey])
, PARTITION ( [ OrderDateKey ] RANGE RIGHT FOR VALUES
                (20170101, 20180101, 20190101, 20200101, 20210101)
            )
)

```

Box 1: HASH

Box 2: OrderDateKey

In most cases, table partitions are created on a date column.

A way to eliminate rollbacks is to use Metadata Only operations like partition switching for data management.

For example, rather than execute a DELETE statement to delete all rows in a table where the order_date was in October of 2001, you could partition your data early. Then you can switch out the partition with data for an empty partition from another table.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-dedicated-sql-pool> Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that the data is distributed evenly across partitions.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 122

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and an Azure Data Lake Storage Gen2 account named Account1.

You plan to access the files in Account1 by using an external table.

You need to create a data source in Pool1 that you can reference when you create the external table.

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.

```
CREATE EXTERNAL DATA SOURCE source1
WITH
    ( LOCATION = 'https://account1.
      [dropdown: blob, dfs, table] core.windows.net'
      [dropdown: PUSHDOWN = ON, TYPE = BLOB_STORAGE, TYPE = HADOOP]
    )
```

Answer:

```
CREATE EXTERNAL DATA SOURCE source1
WITH
    ( LOCATION = 'https://account1.
      [dropdown: blob, dfs, table] core.windows.net'
      [dropdown: PUSHDOWN = ON, TYPE = BLOB_STORAGE, TYPE = HADOOP]
    )
```

Explanation:



Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2 `CREATE EXTERNAL DATA SOURCE YellowTaxi WITH ( LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/ ', TYPE = HADOOP)` Box 2: HADOOP Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a data source in Pool1 that you can reference when you create the external table.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 123

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

Inherit from: ⓘ

Server

Azure defaults

Don't inherit

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

Configure the automatic tuning options ⓘ

OPTION		DESIRED STATE	CURRENT STATE
	FORCE PLAN	ON OFF INHERIT	ON Auto-configured by Azure
	CREATE INDEX	ON OFF INHERIT	ON Auto-configured by Azure
	DROP INDEX	ON OFF INHERIT	ON Forced by user

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

NOTE: Each correct selection is worth one point.

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

Answer:

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.	☒	☐

Explanation:

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.	☐	☐

Box 1: Yes

We see: Tuning option: Create index ON

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

Box 2: No

Box 3: Yes

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

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure SQL database named DB1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 124

You have 25 Azure SQL databases.

You need to implement a centralized database management solution that uses Transact-SQL What should you include in the solution?

- A. elastic jobs
- B. an Azure Automation runbook
- C. Azure Functions
- D. Azure Logic Apps

Answer: (SHOW ANSWER)

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why B is Correct: Azure Automation runbooks execute PowerShell or Python tasks on a schedule, making them useful for Azure SQL maintenance workflows. This fits the case because the requested outcome is: You need to implement a centralized database management solution that uses Transact-SQL What should you include in the solution? The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Elastic Jobs automate T-SQL execution across many Azure SQL databases, especially for multi-tenant or shard-style deployments. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why C is Wrong: Azure Functions can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why D is Wrong: Azure Logic Apps can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

NEW QUESTION: 125

You have an Azure subscription that contains a SQL Server on Azure Virtual Machines instance named SQLVM1. SQLVM1 has the following configurations:

- * Automated patching is enabled.
- * The SQL Server IaaS Agent extension is installed.
- * The Microsoft SQL Server instance on SQLVM1 is managed by using the Azure portal.

You need to automate the deployment of cumulative updates to SQLVM1 by using Azure Update Manager.

The solution must ensure that the SQL Server instance on SQLVM1 can be managed by using the Azure portal.

What should you do first on SQLVM1?

- A.** Install the Azure Monitor Agent.
- B.** Uninstall the SQL Server IaaS Agent extension.
- C.** Install the Log Analytics agent.
- D.** Set Automated patching to Disable

Answer: C (LEAVE A REPLY)

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why C is Correct: Install the Log Analytics agent. is correct because it is the feature whose normal purpose matches the stated requirement. Log Analytics stores logs and metrics in a queryable workspace and supports workbooks, dashboards, alerting, and cross-resource analysis. The scenario wording points to that specific behavior: You need to automate the deployment of cumulative updates to SQLVM1 by using Azure Update Manager.

Why A is Wrong: Install the Azure Monitor Agent. can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why B is Wrong: Uninstall the SQL Server IaaS Agent extension. can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why D is Wrong: Set Automated patching to Disable can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

NEW QUESTION: 126

You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8. You need to deploy a failover cluster instance (FCI) to VM1 and VM2. The solution must eliminate the need for the following:

- * A distributed network name (DNN)
- * A load balancer

What should you do?

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Connect VM1 and VM2 to different subnets on a single virtual network. is correct because it is the feature whose normal purpose matches the stated requirement. Connect VM1 and VM2 to different subnets on a single virtual network. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to deploy a failover cluster instance (FCI) to VM1 and VM2.

Why A is Wrong: Deploy VM1 and VM2 to a single proximity placement group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: Deploy VM1 and VM2 to different proximity placement groups in the same Azure region.

is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: Connect VM1 and VM2 to a single subnet. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 127

You have an Azure subscription.

You create a logical SQL server that hosts four databases Each database will be used by a separate customer.

You need to ensure that each customer can access only its own database. The solution must minimize administrative effort 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 network security group (NSG)
- B. Create a server-level firewall rule
- C. Create a private endpoint
- D. Create a database-level firewall rule.
- E. Deny public access.

Answer: (SHOW ANSWER)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why C and D is Correct: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The scenario asks for: You need to ensure that each customer can access only its own database. That makes Create a private endpoint the option that satisfies the required Azure SQL layer and operational outcome. Create a database-level firewall rule. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to ensure that each customer can access only its own database. Create a database-level firewall rule. satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Create a network security group (NSG) is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that each customer can access only its own database.

Why B is Wrong: Create a server-level firewall rule is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why E is Wrong: Deny public access. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 128

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

Clustered columnstore indexes

Automatic tuning

Change tracking

PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

A. Clustered columnstore indexes

B. PolyBase

C. Change tracking

D. Automatic tuning

Answer: (SHOW ANSWER)

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

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

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: PolyBase is correct because it is the feature whose normal purpose matches the stated requirement. PolyBase is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

Why A is Wrong: Clustered columnstore indexes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Change tracking is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 129

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

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

You need to configure a connection between VM1 and MIL The solution must meet the following requirements:

- * The connection must be encrypted.
- * Network latency must be minimized.

What should you implement?

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

Answer: [\(SHOW ANSWER\)](#)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The scenario asks for: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted. That makes private endpoints the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: virtual network peering is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted.

Why C is Wrong: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: a site-to-site VPN is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted.

NEW QUESTION: 130

You have a single availability set that contains two SQL Server on Azure Virtual Machines instances. The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1. You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:

- * Provide an availability SLA
- * Minimize costs.

What should you create?

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

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: a distributed network name (DNN) resource is correct because it is the feature whose normal purpose matches the stated requirement. a distributed network name (DNN) resource is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to ensure that client applications can connect to FCI1.

Why A is Wrong: a virtual network name (VNN) resource is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: a Basic Azure Load Balancer is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: an Azure Standard Load Balancer is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 131

You have an Azure subscription that contains an Azure SQL database named DB1.


You need to migrate DB1 to a vCore-based purchasing model by using the Azure Command-Line Interface (Azure CLI). The solution must meet the following requirements:

- * Include two high availability replicas.
- * Minimize downtime.

How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.





Explanation:



Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to migrate DB1 to a vCore-based purchasing model by using the Azure Command-Line Interface (Azure CLI).

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

NEW QUESTION: 132

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

Solution: You create additional tempdb files.

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why A is Correct: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: Each question in the series contains a unique solution that might meet the stated goals. The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: No is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 133

You have a Microsoft SQL Server 2017 server.

You need to migrate the server to Azure. The solution must meet the following requirements:

- * Ensure that the latest version of SQL Server is used.
- * Support the SQL Server Agent service.

Minimize administrative effort.

What should you use?

- A.** SQL Server on Azure Virtual Machines
- B.** Azure SQL Database
- C.** an Azure SQL Database elastic pool
- D.** Azure SQL Managed Instance

Answer: A ([LEAVE A REPLY](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why A is Correct: SQL Server on Azure Virtual Machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario.

The scenario asks for: You need to migrate the server to Azure. That makes SQL Server on Azure Virtual Machines the option that satisfies the required Azure SQL layer and operational outcome.

Why B is Wrong: Azure SQL Database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.


Why C is Wrong: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: Azure SQL Managed Instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

NEW QUESTION: 134

You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



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

Operator: Greater than

Aggregation type: Total

Threshold value: 2

Aggregation granularity (Period): 5 minutes

Frequency of evaluation: Every 5 minutes

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

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

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	☐	☐
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.	☐	☒

Explanation:

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.	☐	☒

Box 1: No

Just one failure within the 5-minute interval.

Box 2: No

Just two failures within the 5-minute interval.

Box 3: No

Just two failures within the 5-minute interval.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

NEW QUESTION: 135

Task 10
You need to protect all the databases on sql37006S95 from SQL injection attacks.

Answer:
See the explanation part for the complete Solution
Explanation:

SQL injection attacks are a type of cyberattack that exploit a vulnerability in the application code that interacts with the database. An attacker can inject malicious SQL statements into the user input, such as a form field or a URL parameter, and execute them on the database server, resulting in data theft, corruption, or unauthorized access1.

To protect all the databases on sql37006S95 from SQL injection attacks, you need to follow some best practices for securing your application and database layers. Here are some of the recommended steps:

Use parameterized queries or stored procedures to separate the SQL code from the user input. This will prevent the user input from being interpreted as part of the SQL statement and avoid SQL injection23.

Validate and sanitize the user input before passing it to the database. This will ensure that the input conforms to the expected format and type, and remove any potentially harmful characters or keywords4.

Implement least privilege access for the database users and roles. This will limit the permissions and actions that the application can perform on the database, and reduce the impact of a successful SQL injection attack5.

Enable Advanced Threat Protection for Azure SQL Database. This is a feature that detects and alerts you of anomalous activities and potential threats on your database, such as SQL injection, brute force attacks, or unusual access patterns. You can configure the alert settings and notifications using the Azure portal or PowerShell.

These are some of the steps to protect all the databases on sql37006S95 from SQL injection attacks.

NEW QUESTION: 136

You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled.

You need to create a new database user for app1.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:



Explanation:



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

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a new database user for app1.

Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 137

You have an Azure subscription that contains an Azure SQL database named DB1.

Users report that DB1 has query performance issues.

You need to ensure that you can move specific tables in DB1 to In-Memory OLTP.

How should you complete the T-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

SELECT DATABASEPROPERTYEX(DB_NAME(),

ALTER DATABASE CURRENT SET

IsXTPSupported;
IsMemoryOptimizedElevateToSnapshotEnabled;
IsXTPSupported;
Updateability;
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT
COMPATIBILITY_LEVEL
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT
READ_COMMITTED_SNAPSHOT

= ON;

Answer:

Answer Area

SELECT DATABASEPROPERTYEX(DB_NAME(),

ALTER DATABASE CURRENT SET

IsXTPSupported;
IsMemoryOptimizedElevateToSnapshotEnabled;
IsXTPSupported;
Updateability;
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT
COMPATIBILITY_LEVEL
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT
READ_COMMITTED_SNAPSHOT

= ON;

Explanation:

Answer Area

SELECT DATABASEPROPERTYEX(DB_NAME(),

ALTER DATABASE CURRENT SET

IsXTPSupported;
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT

= ON;

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that you can move specific tables in DB1 to In-Memory OLTP.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 138

You create an Azure SQL managed instance and a job that performs backups.

You need to configure the job to notify a distribution group by email when the job fails. 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions Commands Cmdlets Statements

⋮ Configure SendGrid.

⋮ Configure an alert.

⋮ Configure Database Mail.

⋮ Create an operator.

⋮ Configure a job notification.

Answer Area

Answer:

Actions Commands Cmdlets Statements

⋮ Configure SendGrid.

⋮ Configure an alert.

⋮ Configure Database Mail.

⋮ Create an operator.

⋮ Configure a job notification.

Answer Area

⋮ Configure Database Mail.

⋮ Create an operator.

⋮ Configure a job notification.

Explanation:

Actions Commands Cmdlets Statements

⋮ Configure SendGrid.

⋮ Configure an alert.

Answer Area

1 ⋮ Configure Database Mail.

2 ⋮ Create an operator.

3 ⋮ Configure a job notification.

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to configure the job to notify a distribution group by email when the job fails.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 139

Task 1

In an Azure SQL database named db1, you need to enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.SalesOrderHeader table.

Answer:

See the explanation part for the complete Solution.

Explanation:

To enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.

SalesOrderHeader table in db1, you can use the following Transact-SQL script:

-- Connect to the Azure SQL database named db1

USE db1;

GO

-- Enable page compression on the clustered index

ALTER INDEX PK_SalesOrderHeader_SalesOrderID ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO This script will rebuild the clustered index with page compression, which can reduce the storage space and improve the query performance The script solution consists of three parts:

The first part is USE db1; GO. This part connects to the Azure SQL database named db1, where the SalesLT.

SalesOrderHeader table is located. The GO command separates the batches of Transact-SQL statements and sends them to the server.

The second part is ALTER INDEX PK_SalesOrderHeader_SalesOrderID ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO. This part enables page compression on the clustered index named PK_SalesOrderHeader_SalesOrderID, which is defined on the SalesLT.

SalesOrderHeader table. The ALTER INDEX statement modifies the properties of an existing index.

The REBUILD option rebuilds the index from scratch, which is required to change the compression setting.

The DATA_COMPRESSION = PAGE option specifies that page compression is applied to the index, which means that both row and prefix compression are used. Page compression can reduce the storage space and improve the query performance by compressing the data at the page level. The GO command ends the batch of statements.

The third part is optional, but it can be useful to verify the compression status of the index. It is SELECT name, index_id, data_compression_desc FROM sys.indexes WHERE object_id = OBJECT_ID(' SalesLT.SalesOrderHeader ');. This part queries the sys.indexes catalog view, which contains information about the indexes in the database. The SELECT statement returns the name, index_id, and data_compression_desc columns for the indexes that belong to the SalesLT.SalesOrderHeader table. The OBJECT_ID function returns the object identification number for the table name. The data_compression_desc column shows the compression type of the index, which should be PAGE for the clustered index after the script is executed.

These are the steps of the script solution for enabling page compression on the clustered index of the SalesLT.

SalesOrderHeader table in db1.

NEW QUESTION: 140

You have an instance of SQL Server on Azure Virtual Machines named SQL1. SQL1 contains a database named DB1.

You need to enable Transparent Data Encryption (TDE) for DB1.

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

Objects

a database encryption key in the master database

a master key in DB1

a certificate in DB1

a master key in the master database

a certificate in the master database

a database encryption key in DB1

Answer Area

>

<

<

>

Answer:

Objects

a database encryption key in the master database

a master key in DB1

a certificate in DB1

a master key in the master database

a certificate in the master database

a database encryption key in DB1

Answer Area

>

<

<

>

1

2

3

a master key in the master database

a certificate in the master database

a database encryption key in DB1

Explanation:

Objects

a database encryption key in the master database

a master key in DB1

a certificate in DB1

Answer Area

1

2

3

a master key in the master database

a certificate in the master database

a database encryption key in DB1

Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to enable Transparent Data Encryption (TDE) for DB1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 141

Task 8

You need to increase the maximum storage size of the query store on db1 to 4,000 MB You may need to use SQL Server Management Studio and the Azure portal.

Answer:

See the explanation part for the complete Solution.

Explanation:

Connect to db1 in SQL Server Management Studio and run:

```
ALTER DATABASE [db1]
SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 4000);
```

Microsoft documents that Query Store maximum storage size is changed with:

```
ALTER DATABASE < database_name >
SET QUERY_STORE (MAX_STORAGE_SIZE_MB = < size > );
```

For Azure SQL Database, the maximum allowed MAX_STORAGE_SIZE_MB value is 10,240 MB, so 4,000 MB is valid.

Method 1 - SSMS / T-SQL Method

This is the cleanest method for the simulation.

Step 1: Open SQL Server Management Studio

Open SQL Server Management Studio.

Connect to the Azure SQL logical server that hosts db1.

Use SQL authentication or Microsoft Entra authentication.

In Connection Properties, select:

db1

You can also connect to master, but for this task it is cleaner to open a query window in db1.

Step 2: Run the Query Store configuration command

Open a new query window and run:

```
ALTER DATABASE [db1]
SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 4000);
```

This changes the Query Store maximum allocated storage size to 4,000 MB.

Step 3: Verify the setting

Run this query in db1:

```
SELECT
actual_state_desc,
desired_state_desc,
current_storage_size_mb,
max_storage_size_mb
FROM sys.database_query_store_options;
```

Expected result:

```
max_storage_size_mb = 4000
```

Microsoft documents sys.database_query_store_options as the catalog view that returns Query Store configuration, including max_storage_size_mb.

Method 2 - SSMS Graphical Method

Use this if the simulation expects a GUI action.

Step 1: Connect to the server

Open SQL Server Management Studio.

Connect to the Azure SQL logical server.

Expand Databases.

Right-click db1.

Select Properties.

Step 2: Open Query Store settings

In the Database Properties window, select Query Store.

Find the setting:

Max Size (MB)

or:

Maximum Storage Size (MB)

Change the value to:

4000

Select OK.

Step 3: Verify

Run:

```
SELECT max_storage_size_mb
FROM sys.database_query_store_options;
```

Expected value:

4000

Azure Portal Clarification

The Azure portal is not the best tool for this specific Query Store configuration. Query Store size is a database engine setting, not a normal Azure resource setting like backup retention, locks, alerts, or private endpoints.

Use SSMS/T-SQL for this task.

Important Notes

Query Store must remain enabled

Do not disable Query Store. Azure SQL Database uses Query Store heavily, and Microsoft notes that Query Store is enabled by default for Azure SQL Database. The task only asks to increase the maximum storage size, not to change capture mode, cleanup policy, or operation mode.

4,000 MB is valid

Azure SQL Database allows Query Store MAX_STORAGE_SIZE_MB up to 10,240 MB, so this setting is within the allowed range.

Final Exam-Lab Action

Run this against db1:

```
ALTER DATABASE [db1]
SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 4000);
```

Then verify with:

```
SELECT max_storage_size_mb
FROM sys.database_query_store_options;
```

The task is complete when the result shows:

4000

NEW QUESTION: 142

A company plans to use Apache Spark analytics to analyze intrusion detection data.

You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts.

What should you recommend?

- A.** Azure Data Lake Storage
- B.** Azure Databricks
- C.** Azure HDInsight
- D.** Azure Data Factory

Answer: (SHOW ANSWER)

Azure HDInsight offers pre-made, monitoring dashboards in the form of solutions that can be used to monitor the workloads running on your clusters. There are solutions for Apache Spark, Hadoop, Apache Kafka, live long and process (LLAP), Apache HBase, and Apache Storm available in the Azure Marketplace.

Note: With Azure HDInsight you can set up Azure Monitor alerts that will trigger when the value of a metric or the results of a query meet certain conditions. You can condition on a query returning a record with a value that is greater than or less than a certain threshold, or even on the number of results returned by a query. For example, you could create an alert to send an email if a Spark job fails or if a Kafka disk usage becomes over 90 percent full.

Reference:

<https://azure.microsoft.com/en-us/blog/monitoring-on-azure-hdinsight-part-4-workload-metrics-and-logs/>

NEW QUESTION: 143

You have an Azure SQL Database managed instance.

The instance starts experiencing performance issues.

You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.

What should you use?

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

Answer: (SHOW ANSWER)

Reference:

[https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the- query-store?view=sql-server-ver15](https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store?view=sql-server-ver15) Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: dynamic management views matches the expected DP-300 administration action. dynamic management views is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to identify which query is causing the issue and retrieve the execution plan for the query.

Why A is Wrong: the Azure portal is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Extended Events is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 144

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1.

The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

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

Actions

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.

Answer Area

Answer:

Actions

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.

Answer Area

Set the Active Directory Admin for AzSQL1.

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Explanation:

Set the Active Directory Admin for AzSQL1.

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Step 1: Set up the Active Directory Admin for AzSQL1.

Step 2: Connect to DB1 by using the server administrator.

Sign into your managed instance with an Azure AD login granted with the sysadmin role.

Step 3: Create a user by using the FROM EXTERNAL PROVIDER clause.

FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:

```
CREATE USER [AAD_principal] FROM LOGIN [Azure AD login]
```

Reference:

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

NEW QUESTION: 145

You have an Azure SQL database named DB1.

You need to display the estimated execution plan of a query by using the query editor in the Azure portal.

What should you do first?

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

Answer: (SHOW ANSWER)

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-showplan-all-transact-sql?view=sql-server-ver15> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why A is Correct: Run the set showplan_all Transact-SQL statement. matches the expected DP-300 administration action. Run the set showplan_all Transact-SQL statement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to display the estimated execution plan of a query by using the query editor in the Azure portal.

Why B is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Run the set forceplan Transact-SQL statement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 146

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

- A.

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

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

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

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

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why C is Correct: Option C is correct because it is the feature whose normal purpose matches the stated requirement. Option C is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: You need to change the password for the private key.

Why A is Wrong: Option A is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why B is Wrong: Option B is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: Option D is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

NEW QUESTION: 147

You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals.

What should you include in the recommendation?

- A.

For each customer, grant the customer ddl_admin to the existing schema.
- B.

For each customer, create an additional schema and grant the customer ddl_admin to the new schema.
- C.

For each customer, create an additional schema and grant the customer db_writerto the new schema.
- D.

For each customer, grant the customer db_writerto the existing schema.

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: For each customer, create an additional schema and grant the customer ddl_admin to the new schema. matches the expected DP-300 administration action. For each customer, create an additional schema and grant the customer ddl_admin to the new schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The question is not asking for a general Azure capability; it is asking for the feature that produces this result:

You need to recommend a solution to ensure that the customers can create the database objects.

Why A is Wrong: For each customer, grant the customer ddl_admin to the existing schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: For each customer, create an additional schema and grant the customer db_writerto the new schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross- region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: For each customer, grant the customer db_writerto the existing schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 148

You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.

You need to create a surrogate key for the table. The solution must provide the fastest query performance.

What should you use for the surrogate key?

A. an IDENTITY column

B. a GUID column

C. a sequence object

Answer: (SHOW ANSWER)

Dedicated SQL pool supports many, but not all, of the table features offered by other databases.

Surrogate keys are not supported. Implement it with an Identity column.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesoverview> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: an IDENTITY column is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to create a surrogate key for the table. The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: a GUID column is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: a sequence object is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

NEW QUESTION: 149

You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image.

You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance.

What should you do first?

A. From the SQL Server Installation Center on SQLVM1. run the Edition Upgrade wizard.

B. From SQLVM1, uninstall the SQL Server instance.

C. From the SQL Server Installation Center on SQLVM1. run the Repair wizard.

D. From the Azure portal, reconfigure SQLVM1.

Answer: (SHOW ANSWER)

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: From SQLVM1, uninstall the SQL Server instance. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario asks for: You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. That makes From SQLVM1, uninstall the SQL Server instance. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: From the SQL Server Installation Center on SQLVM1. run the Edition Upgrade wizard. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: From the SQL Server Installation Center on SQLVM1. run the Repair wizard. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: From the Azure portal, reconfigure SQLVM1. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

NEW QUESTION: 150

You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.

You need to recommend a deployment solution that meets the following requirements:

- * Provides a Service Level Agreement (SLA) of at least 99.95%
- * Replicates databases in the same group synchronously
- * Minimizes the latency of database writes

What should you recommend?

- A.** Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group.
- B.** Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group.
- C.** Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group.
- D.** Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group.

Answer: (SHOW ANSWER)

To get VMs as close as possible, achieving the lowest possible latency, you should deploy them within a proximity placement group. <https://learn.microsoft.com/en-us/azure/virtual-machines/co-location> Basic Concept:

This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why A is Correct: Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group. is correct because it is the feature whose normal purpose matches the stated requirement. Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to recommend a deployment solution that meets the following requirements: * Provides a Service Level Agreement (SLA) of at least 99.95% * Replicates databases in the same group synchronously * Minimizes the latency of database writes What should you r

Why B is Wrong: Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 151

You have an Azure AD tenant and a logical Microsoft SQL server named SQL1 that hosts several Azure SQL databases.

You plan to assign Azure AD users permissions to the databases automatically by using Azure Automation.

You need to create the required Automation accounts.

Which two accounts should you create? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A.** From the Azure Active Directory admin center, create a service principal.
- B.** From the Azure Active Directory admin center, create a user-assigned managed identity for SQL1.
- C.** On SQL1. create a SQL user in the databases.
- D.** On SQL1, create a SQL login.
- E.** From the Azure Active Directory admin center, create an external identity.

Answer: (SHOW ANSWER)

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why A and C is Correct: Active Directory or Microsoft Entra integration centralizes identity and group membership so database access can be managed through directory accounts. In this scenario, the important constraint is: You need to create the required Automation accounts. From the Azure Active Directory admin center, create a service principal. satisfies that constraint without adding an unrelated service or manual process. On SQL1. create a SQL user in the databases. matches the expected DP-300 administration action.

On SQL1. create a SQL user in the databases. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to create the required Automation accounts.

Why B is Wrong: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: On SQL1, create a SQL login. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why E is Wrong: Active Directory or Microsoft Entra integration centralizes identity and group membership so database access can be managed through directory accounts. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 152

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

DBCC CHECK [DB1],

REPAIR_REBUILD

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

PHYSICAL_ONLY;

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

Answer:

DBCC CHECK [DB1],

REPAIR_REBUILD

NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

PHYSICAL_ONLY;

EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

Explanation:
Answer Area

DBCC CHECK [DB1], REPAIR_REBUILD) WITH PHYSICAL_ONLY

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to identify and repair any physical or logical corruption in DB1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

NEW QUESTION: 153

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.

You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.

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

Actions	Answer Area
Encrypt the Salary column by using the randomized encryption type.	
Create a column encryption key.	
Enable Transparent Data Encryption (TDE).	
Encrypt the Salary column by using the deterministic encryption type.	
Apply a dynamic data mask to the Salary column.	
Create a column master key.	

Answer:

Actions

Encrypt the Salary column by using the randomized encryption type.

Create a column encryption key.

Enable Transparent Data Encryption (TDE).

Encrypt the Salary column by using the deterministic encryption type.

Apply a dynamic data mask to the Salary column.

Create a column master key.

Answer Area

Create a column master key.

Create a column encryption key.

Encrypt the Salary column by using the randomized encryption type.

Explanation:

Create a column master key.

Create a column encryption key.

Encrypt the Salary column by using the randomized encryption type.

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION: 154

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

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

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use 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> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why A is Correct: Yes is correct because it is the feature whose normal purpose matches the stated requirement. Yes is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: Each question in the series contains a unique solution that might meet the stated goals.

Why B is Wrong: No is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

NEW QUESTION: 155

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL.

Which switch should you use to switch between languages?

A. \[< language >]

B. % < language >

C. \[< language >]

D. @ < language >

Answer: ([SHOW ANSWER](#))

Explanation/Reference:

You can override the default language by specifying the language magic command % < language > at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use> Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: % < language > is correct because it is the feature whose normal purpose matches the stated requirement. % < language > is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: Which switch should you use to switch between languages?

Why A is Wrong: \[< language >] is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why C is Wrong: \[< language >] is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: @ < language > is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

NEW QUESTION: 156

You have an Azure SQL database named db1 on a server named server1.

You use Query Performance Insight to monitor db1.

You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Configuration setting:

Microsoft

DATA_FLUSH_INTERVAL_SECONDS

INTERVAL_LENGTH_MINUTES

MAX_PLANS_PER_QUERY

QUERY_CAPTURE_MODE

Value:

1

60

CUSTOM

ON

Answer:

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS

INTERVAL_LENGTH_MINUTES

MAX_PLANS_PER_QUERY

QUERY_CAPTURE_MODE

Value:

1

60

CUSTOM

ON

Explanation:

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS

INTERVAL_LENGTH_MINUTES

MAX_PLANS_PER_QUERY

QUERY_CAPTURE_MODE

Value:

1

60

CUSTOM

ON

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

NEW QUESTION: 157

You plan to deploy Instance1 by using the following script.

```
{
  "type": "Microsoft.Sql/managedInstances",
  "apiVersion": "2019-06-01-preview",
  "name": "[parameters('instanceName')]",
  "location": "[resourceGroup().location]",
  "sku": {
    "name": "[parameters('skuName')]",
    "tier": "[parameters('skuEdition')]"
  },
  "dependsOn": [
    "Microsoft.Resources/deployments/BuildMInetworking"
  ],
  "properties": {
    "administratorLogin": "[parameters('adminLogin')]",
    "administratorLoginPassword": "[parameters('adminPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',parameters('netName'),parameters('subnetName'))]"
  }
}
```

You need to specify the licenseType and storagedundancy parameters. The deployment must meet the availability requirements and the business requirements for DB1 and DB2. To what should you set each parameter? To answer, select the appropriate options in the answer area.

Answer Area

licenseType:

storageRedundancy:

Microsoft

Answer:

Answer Area

licenseType:

storageRedundancy:

Microsoft

Explanation:

Answer Area

licenseType:

storageRedundancy:

Microsoft

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to specify the licenseType and storagedundancy parameters.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

NEW QUESTION: 158

You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.

The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb on sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Answer: (SHOW ANSWER)

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting.

You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb-contention>

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-use-diagnostics-log> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why D is Correct: Run dynamic management view-based queries is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to identify whether the performance issue relates to the use of tempdb on sqlsrv1. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why B is Wrong: Review information provided by SQL Server Profiler-based traces is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Query Performance Insight surfaces resource-consuming queries and database workload patterns in Azure SQL for performance investigation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 159

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> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D is Correct: Increase the number of streaming units (SUs). is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to ensure that the job can handle all the events. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Remove any named consumer groups from the connection and use \$default. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why B is Wrong: Change the compatibility level of the Stream Analytics job. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Create an additional output stream for the existing input stream. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

NEW QUESTION: 160

You need to recommend a process to automate the management of DB3. The solution must meet the management requirements. What should be the first step of the process?

A. Configure Microsoft Entra authentication for the logical server that hosts DB3.

- B.** Create a database that has database-scoped credentials.
- C.** Configure a private endpoint for connectivity to DB3.
- D.** Create data base-scoped credentials in DB3.

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why C is Correct: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The scenario asks for: You need to recommend a process to automate the management of DB3. That makes Configure a private endpoint for connectivity to DB3. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Configure Microsoft Entra authentication for the logical server that hosts DB3. is a security- related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: Create a database that has database-scoped credentials. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to recommend a process to automate the management of DB3.

Why D is Wrong: Create data base-scoped credentials in DB3. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

NEW QUESTION: 161

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 User1 creates.
- D.** Grant select on sys.databases to public in the master database.

Answer: ([SHOW ANSWER](#))

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage> Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why A is Correct: Create User1 in the master database. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases. Create User1 in the master database. satisfies that constraint without adding an unrelated service or manual process.

Why B is Wrong: Assign User1 the db_datareader role for the master database. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why C is Wrong: Assign User1 the db_datareader role for the databases that User1 creates. is a security- related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: Grant select on sys.databases to public in the master database. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.

NEW QUESTION: 162

You configure backups for an Azure SQL database as shown in the following exhibit.

Point-in-time-restore

Specify how long you want to keep your point-in-time backups. [Learn more](#)

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

14

Long-term retention

Specify how long you want to keep your long-term retention backups. You may choose to keep yearly backups for up to 10 years. [Learn more](#)

Weekly LTR Backups

Keep weekly backups for:

52

Week(s)▼

Monthly LTR Backups

Keep the first backup of each month for:

60

Week(s)▼

Yearly LTR Backups

Keep an annual backup for:

10

Year(s)▼

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

Week 52

▼

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

NOTE: Each correct selection is worth one point.

Answer Area

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

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

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

65 backup copies
1 backup copy
52 backup copies
64 backup copies
65 backup copies

Answer:

Answer Area

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

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

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

65 backup copies
1 backup copy
52 backup copies
64 backup copies
65 backup copies

Explanation:

Answer Area

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

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

point-time restore (PITR) backup.

65 backup copies

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You configure backups for an Azure SQL database as shown in the following exhibit.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

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. Transact-SQL

C. The Azure portal

D. Azure PowerShell

Answer: ([SHOW ANSWER](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: The Azure portal is correct because it is the feature whose normal purpose matches the stated requirement. The Azure portal is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to perform a point-in-time restore of DB1 to SQLMI2.

Why A is Wrong: Azure CLI is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: Transact-SQL is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: Azure PowerShell is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

NEW QUESTION: 164

Task 5

You need to generate an email alert for db1 if the average CPU percentage utilization is greater than 50 percent for five minutes sampled at one-minute intervals. The alert must be sent to admin@contoso.com.

You may need to use SQL Server Management Studio and the Azure portal.

Answer:

See the explanation part for the complete Solution.

Explanation:

Create an Azure Monitor metric alert rule on database db1 using the metric:

CPU percentage

Configure the condition as:

Aggregation: Average

Operator: Greater than

Threshold: 50

Aggregation granularity / Period: 1 minute

Frequency of evaluation: 1 minute

Evaluation period / Lookback window: 5 minutes

Action group email: admin@contoso.com

Azure SQL Database exposes CPU percentage as a platform metric, and Azure Monitor metric alerts can send notifications through action groups such as email. Microsoft describes SQL Database alerts as metric-based alerts that can send email when metrics such as CPU usage reach a defined threshold.

Azure Portal Method - Recommended for Simulation

Step 1: Open the db1 Azure SQL database

Sign in to the Azure portal.

Search for SQL databases.

Open the database named:

db1

Do not open the SQL logical server unless the alert needs to apply to all databases. This task is specifically for db1, so the alert scope must be the db1 database resource.

Step 2: Create a new alert rule

From the db1 database page:

In the left menu, select Alerts.

Select Create.

Select Alert rule.

Microsoft's Azure Monitor workflow allows you to create an alert rule directly from the target resource.

When you create it from a resource, the resource is automatically set as the alert scope.

Step 3: Confirm the alert scope

On the alert rule page, confirm the scope is the Azure SQL database:

db1

The resource type should be similar to:

SQL database

Microsoft.Sql/servers/databases

If the scope is the SQL server instead of the database, remove it and select the db1 database resource.

Step 4: Add the alert condition

Under Condition, select Add condition.

Choose the metric:

CPU percentage

Sometimes it appears as:

CPU percent

or metric name:

cpu_percent

For Azure SQL Database, CPU percentage represents CPU consumption toward the database workload limit, expressed as a percentage.

Step 5: Configure the signal logic

Configure the alert logic exactly like this:

Setting

Value

Threshold type

Static

Aggregation type

Average

Operator

Greater than

Threshold value

50

Unit

Percent

Aggregation granularity / Period

1 minute

Frequency of evaluation

1 minute

Evaluation period / Lookback window

5 minutes

This means Azure Monitor evaluates the CPU metric every minute, using one-minute metric samples, and fires the alert only when the average CPU percentage is greater than 50% across the five-minute evaluation window.

Be precise here. The task says:

average CPU percentage utilization is greater than 50 percent for five minutes sampled at one-minute intervals So the correct choices are:

Average

Greater than 50

Every 1 minute

Over the last 5 minutes

Azure Monitor metric alert rules combine the monitored resource, the metric condition, and action groups that run when the condition is met.

Step 6: Create the email action group

Under Actions, select:

Create action group

Configure the basics:

Setting

Value

Subscription

Use the current subscription

Resource group

Use the lab resource group

Action group name

AG-db1-CPU-Email

Display name

db1CPU

Then go to Notifications.

Add a notification:

Setting

Value

Notification type

Email/SMS message/Push/Voice

Name

EmailAdmin

Email

admin@contoso.com

Select OK, then Review + create, then Create.

Action groups define the notification or automation action that runs when an alert fires. Microsoft documents email as a supported action group notification type.

Step 7: Configure alert rule details

Under Details, configure:

Setting

Value

Severity

2 or 3

Alert rule name

db1 CPU greater than 50 percent

Description

Alert when average CPU percentage for db1 is greater than 50 percent for 5 minutes.

Enable upon creation

Yes

Severity is usually not specified by the task, so any reasonable severity is acceptable. In an exam lab, I would use Severity 2 for CPU performance impact.

Step 8: Review and create

Select Review + create.

Confirm the condition says something equivalent to:

Whenever the average CPU percentage is greater than 50

Confirm the evaluation settings are:

Check every 1 minute

Lookback period 5 minutes

Confirm the action group sends email to:

admin@contoso.com

Select Create.

Verification

After creation:

Open db1.

Go to Alerts.

Select Alert rules.

Confirm the alert rule exists and is enabled.

Open the rule and verify:

Scope: db1

Signal: CPU percentage

Aggregation: Average

Operator: Greater than

Threshold: 50

Evaluation frequency: 1 minute

Window size: 5 minutes

Action group: admin@contoso.com

Azure CLI Method

Use this only if the simulation gives you Cloud Shell.

First get the database resource ID:

az sql db show \

```
--resource-group < resource-group-name > \  
--server < sql-server-name > \  
--name db1 \  
--query id \  
--output tsv  
Create the action group:  
az monitor action-group create \  
--resource-group < resource-group-name > \  
--name AG-db1-CPU-Email \  
--short-name db1CPU \  
--action email EmailAdmin admin@contoso.com
```

```
Create the metric alert:  
az monitor metrics alert create \  
--name " db1 CPU greater than 50 percent " \  
--resource-group < resource-group-name > \  
--scopes < db1-resource-id > \  
--condition " avg cpu_percent > 50 " \  
--window-size 5m \  
--evaluation-frequency 1m \  
--action AG-db1-CPU-Email \  
--description " Alert when average CPU percentage for db1 is greater than 50 percent for 5 minutes. " The metric name commonly used for Azure SQL Database CPU percentage in CLI/ARM contexts is:  
cpu_percent
```

SSMS Clarification

SSMS is not the correct tool for this task.
Do not configure Database Mail. Azure SQL Database does not use SQL Server Agent/Database Mail in the same way as SQL Server on a VM or SQL Managed Instance. This requirement is an Azure Monitor metric alert requirement.
Correct tool:

Azure portal > db1 > Alerts > Create alert rule

Wrong tool:

SSMS Database Mail

Final Exam-Lab Configuration

Use this exact configuration:

Resource: db1

Alert type: Metric alert

Metric: CPU percentage

Aggregation: Average

Operator: Greater than

Threshold: 50

Aggregation granularity: 1 minute

Evaluation frequency: 1 minute

Evaluation period/window: 5 minutes

Action group notification: Email

Email recipient: admin@contoso.com

That completes the task.

NEW QUESTION: 165

You have an Azure subscription that contains an Azure SQL database named SQLDB1 in the Hyperscale service tier. SQLDB1 has zone-redundant storage (ZRS) enabled. You need to create a copy of SQLDB1 in the paired Azure region. The solution must minimize administrative effort. What should you do first?

A. Create a logical server for Azure SQL Database.

B. Modify the Backup storage redundancy setting for SQLDB1.

C. Create another Azure SQL database in the Hyperscale service tier.

D. Export SQLDB1.

Answer: (SHOW ANSWER)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why A is Correct: Create a logical server for Azure SQL Database. matches the expected DP-300 administration action. Create a logical server for Azure SQL Database. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to create a copy of SQLDB1 in the paired Azure region.

Why B is Wrong: Modify the Backup storage redundancy setting for SQLDB1. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why C is Wrong: Hyperscale is designed for very large databases and fast scale-out storage architecture, not merely for intermittent CPU bursts. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: Export SQLDB1. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

NEW QUESTION: 166

You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines that supports Write Accelerator.

Which virtual machine series should you use?

A. H-series

B. G-series

C. M-series

D. E-series

Answer: (SHOW ANSWER)

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: M-series matches the expected DP-300 administration action. M-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which virtual machine series should you use?

Why A is Wrong: H-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why B is Wrong: G-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why D is Wrong: E-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (503 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 167

You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil. You plan to perform an online migration of db1 to MM by using Azure Database Migration Service. You need to create the backups for the migration. The solution must minimize the number of backup files created. Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Backup type:

To store the backups:

Answer:

Answer Area

Backup type:

To store the backups:

Explanation:

Answer Area

Backup type: Full only

To store the backups: Append each backup to a separate file.

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create the backups for the migration.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

NEW QUESTION: 168

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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds>

Valid DP-300 Dumps shared by EduDump.com for Helping Passing DP-300 Exam! EduDump.com now offer the **newest DP-300 exam dumps**, the EduDump.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com DP-300 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/DP-300/premium/> (**503** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)