

Microsoft.DP-300.v2026-07-31.q244

Exam Code:	DP-300
Exam Name:	Administering Relational Databases on Microsoft Azure
Certification Provider:	Microsoft
Free Question Number:	244
Version:	v2026-07-31
# of views:	177
# of Questions views:	2440
https://www.freepdfdumps.com/Microsoft.DP-300.v2026-07-31.q244.html	

NEW QUESTION: 1

You have an Azure subscription that contains an Azure SQL database. The database contains a table named table1.
You execute the following Transact-SQL statements.

```
CREATE CLUSTERED INDEX PK_index1 ON [dbo].[table1] ([Column1])
CREATE NONCLUSTERED INDEX [NCI_index2] ON [dbo].[table1] ([Column1]) ASC
```

You need to reduce the time it takes to perform analytic queries on the database.
Which configuration should you enable?

- A. ROW_MODE_MEMORY_GRANT_FEEDBACK
- B. BATCH_MODE_MEMORY_GRANT_FEEDBACK
- C. BATCH_MODE_ADAPTIVE_JOINS
- D. BATCH_MODE_ON_ROWSTORE

Answer: (SHOW ANSWER)

This feature, Batch Mode on Rowstore, allows all operators enabled for batch mode to operate as such in SQL Server.
What does this mean? It means that query operations can process data faster, more efficiently and mimic what makes columnstore so fast. Instead of reading row by row (row store mode) it can read in chunks i.e. batches of 900 rows at a time. The performance impact of this can be tremendous which effectively uses CPU more efficiently.
<https://sqlespresso.com/2020/06/16/what-is-batch-mode-on-rowstore-in-sql-server/>

NEW QUESTION: 2

Hotspot Question
You configure a long-term retention policy for an Azure SQL database as shown in the exhibit.

Configure policies



SQL server

Point in Time Restore Configuration

Configure PiTR backup retention Days

Long-term Retention Configurations

☒ Weekly LTR Backups

How long would you like weekly backups to be kept?

6 Week(s)

☒ Monthly LTR Backups

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

12 Month(s)

☒ Yearly LTR Backups

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

Week 2

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

10 Year(s)

The first weekly backup occurred on January 4, 2020. The dates for the first 10 weekly backups are:

January 4, 2020
January 11, 2020
January 18, 2020
January 25, 2020
February 1, 2020
February 8, 2020
February 15, 2020
February 22, 2020

February 29, 2020

March 7, 2020

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

NOTE: Each correct selection is worth one point.



The backup saved to long-term retention on January 4, 2020, will be retained for

6 weeks

12 months

10 years

The backup saved to long-term retention on January 11, 2020 will be retained for

6 weeks

12 months

10 years

Answer:



The backup saved to long-term retention on January 4, 2020, will be retained for

6 weeks

12 months

10 years

The backup saved to long-term retention on January 11, 2020 will be retained for

6 weeks

12 months

10 years

Explanation:

Box 1: 12 months

The January 4, 2020 backup is the first backup of the month, and it will keep for 12 months.

Box 2: 10 years.

The January 11, 2020 backup is the Week 2 backup. This backup will be the yearly backup, and it will be retained for 10 years.

Note: Long-term backup retention (LTR) leverages the full database backups that are automatically created to enable point in time restore (PITR). If an LTR policy is configured, these backups are copied to different blobs for long-term storage. The copy is a background job that has no performance impact on the database workload. The LTR policy for each database in SQL Database can also specify how frequently the LTR backups are created.

To enable LTR, you can define a policy using a combination of four parameters: weekly backup retention (W), monthly backup retention (M), yearly backup retention (Y), and week of year (WeekOfYear). If you specify W, one backup every week will be copied to the long-term storage.

If you specify M, the first backup of each month will be copied to the long-term storage. If you specify Y, one backup during the week specified by WeekOfYear will be copied to the long-term storage. If the specified WeekOfYear is in the past when the policy is configured, the first LTR backup will be created in the following year. Each backup will be kept in the long-term storage according to the policy parameters that are configured when the LTR backup is created.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/long-term-retention-overview?view=azuresql>

NEW QUESTION: 3

You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.

You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable query optimizer fixes.
- B. Enable Resource Governor.
- C. Enable parameter sniffing.
- D. Create a workload group.
- E. Configure In-memory OLTP.
- F. Run the Database Engine Tuning Advisor.
- G. Reduce the Max Degree of Parallelism value.

Answer: (SHOW ANSWER)

SQL Server Resource Governor is a feature that you can use to manage SQL Server workload and system resource consumption. Resource Governor enables you to specify limits on the amount of CPU, physical I/O, and memory that incoming application requests can use.

The following concept is fundamental to understanding and using Resource Governor:

* Workload groups. A workload group serves as a container for session requests that have similar classification criteria. A workload allows for aggregate monitoring of the sessions, and defines policies for the sessions. Each workload group is in a resource pool. Two workload groups (internal and default) are created and mapped to their corresponding resource pools when SQL Server is installed. Resource Governor also supports user-defined workload groups.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15>

NEW QUESTION: 4

You have four Azure subscriptions. Each subscription contains multiple Azure SQL databases.

You need to update the column and index statistics for the databases.

What should you use?

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

Answer: A (LEAVE A REPLY)

You can create a runbook for index maintenance in an Azure SQL database.

Reference:

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

NEW QUESTION: 5

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.

Drag and Drop Question

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:



Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

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

1. Create and populate Azure AD.
2. Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
3. Create an Azure Active Directory administrator. (Step 1)
4. Configure your client computers.
5. Create contained database users in your database mapped to Azure AD identities. (Step 2)
6. 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: 6

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: B (LEAVE A REPLY)

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).

Incorrect Answers:

A: Locally redundant storage (LRS) copies your data synchronously three times within a single physical location in the primary region. LRS is the least expensive replication option, but is not recommended for applications requiring high availability.

C: Zone-redundant storage (ZRS) copies your data synchronously across three Azure availability zones in the primary region.

D: 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.

Reference:

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

NEW QUESTION: 7

Hotspot Question

You have an Azure SQL database named DB1.

You need to configure DB1 to use the General Purpose service tier.

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

ALTER DATABASE

ALTER RESOURCE GOVERNOR

ALTER SERVER CONFIGURATION

ALTER SERVER ROLE

[DB1] MODIFY (

EDITION

NAME


RECONFIGURE

SERVICE_OBJECTIVE

= 'GeneralPurpose');

Microsoft

Answer:



Explanation:

Box 1: ALTER DATABASE

To configure an Azure SQL Database to use the General Purpose service tier using Transact- SQL, you use the EDITION = 'GeneralPurpose' argument inside a CREATE DATABASE or ALTER DATABASE statement.

Box 2: EDITION

For an Existing Database

If you are modifying or migrating an existing database to the General Purpose tier, use the ALTER DATABASE statement:

ALTER DATABASE MyDatabaseName

MODIFY (EDITION = 'GeneralPurpose', SERVICE_OBJECTIVE = 'GP_Gen5_2'); Reference:

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

NEW QUESTION: 8

You have an Azure SQL database.

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

You run the select * from sys.database_scoped_configurations Transact-SQL command and receive the results shown in the following table.

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

You need relieve the memory pressure.

What should you configure?

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

Answer: C (LEAVE A REPLY)

The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/optimize-for-ad-hoc-workloads-server-configuration-option?view=sql-server-ver15>

NEW QUESTION: 9

You have two Azure virtual machines named Server1 and Server2 that run Windows Server 2022 and are joined to an Active Directory Domain Services (AD DS) domain named contoso.com.

Both virtual machines have a default instance of Microsoft SQL Server 2019 installed. Server1 is configured as a master server, and Server2 is configured as a target server.

On Server1, you create a proxy account named contoso\sqlproxy.

You need to ensure that the SQL Server Agent job steps can be downloaded from Server1 and run on Server2.

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

NOTE: Each correct selection is worth one point.

- A. On Server2, set the AllowDownloadedJobsToMatchProxyName registry entry to 1.
- B. On Server2, grant the contoso\sqlproxy account the Impersonate a client after authentication user right.
- C. On Server2, create a proxy account.
- D. On Server1, set the AllowDownloadedJobsToMatchProxyName registry entry to 1.
- E. On Server2, grant the contoso\sqlproxy account the Access this computer from the network user right.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 10

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.

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.

Seeing a Count of All Active SQL Server Wait Types.

Sometimes we're trying to diagnose a problem and we want to know if we're seeing a large number of wait types occurring. We can do this using sys.dm_exec_requests because the current wait type being experienced is presented. Therefore, if we filter out any background or sleeping tasks, we can get a picture of what the waits are for active requests and we can also see if we have a problem. Here's the query:

```
SELECT COALESCE(wait_type, 'None') AS wait_type, COUNT(*) AS Total
FROM sys.dm_exec_requests -
WHERE NOT status IN ('Background', 'Sleeping')
GROUP BY wait_type
ORDER BY Total DESC;
```

Here is an example of the query output:



	wait_type	Total
1	LCK_M_S	2
2	None	1

We see that we have two LCK_M_S wait types. This is the wait type we get when we have requests waiting on obtaining a shared lock. We can then query along with sys.dm_tran_locks to determine what types of locks these active requests were trying to obtain:

```
SELECT L.request_session_id, L.resource_type,L.resource_subtype, L.request_mode,
L.request_type
FROM sys.dm_tran_locks AS L
JOIN sys.dm_exec_requests AS DER
ON L.request_session_id = DER.session_id
WHERE DER.wait_type = 'LCK_M_S';
```

Incorrect:

Not A: Azure SQL database does not have sys.dm_pdw_nodes_tran_locks.

Not E: Azure SQL database does not have sys.dm_pdw_nodes_os_wait_stats.

Reference:

<https://www.mssqltips.com/sqlservertip/5521/understanding-and-using-sql-server-sysdmexecrequests/>

NEW QUESTION: 11

Case Study 4 - A.Datum

Overview

ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.

Existing Environment

ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.

SALESDB collects data from the stores and the website.

DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.

REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

Requirements

Planned Changes

ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:

Migrate SALESDB and REPORTINGDB to an Azure SQL database.

Migrate DOCDB to Azure Cosmos DB.

The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.

As they arrive, all the sales documents in JSON format must be transformed into one consistent format.

Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.

Technical Requirements

The new Azure data infrastructure must meet the following technical requirements:

Data in SALESDB must be encrypted by using Transparent Data Encryption (TDE). The encryption must use your own key.

SALESDB must be restorable to any given minute within the past three weeks.

Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

Missing indexes must be created automatically for REPORTINGDB.

Disk IO, CPU, and memory usage must be monitored for SALESDB.

Which counter should you monitor for real-time processing to meet the technical requirements?

A. SU% Utilization

B. CPU% utilization

C. Concurrent users

D. Data Conversion Errors

Answer: A (LEAVE A REPLY)

Scenario:

- Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

- The sales data including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics.

Streaming Units (SUs) represents the computing resources that are allocated to execute a Stream Analytics job. The higher the number of SUs, the more CPU and memory resources are allocated for your job.

This capacity lets you focus on the query logic and abstracts the need to manage the hardware to run your Stream Analytics job in a timely manner.

References:


<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-streaming-unit-consumption>

NEW QUESTION: 12

Hotspot Question

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

Yes

No

You will reduce the I/O usage and the query execution time if you force the query plan.

☐ Yes ☐ No

You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.

☐ Yes ☐ No

You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.

☐ Yes ☐ No

Answer:

Answer Area



Statements

You will reduce the I/O usage and the query execution time if you force the query plan.

Yes

☐

No

☒

You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.

☐☒

You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.

☒☐

Explanation:

Box 1: No

There is only one query plan available. Force has no effect.

Box 2: No

Adding an index will not increase IO usage.

Box 3: Yes

The performance would improve.

Reference:

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

NEW QUESTION: 13

Drag and Drop Question

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		Answer Area		
a database encryption key in the master database	> <	1		↑ ↓
a master key in the master database		2		
a certificate in DB1		3		
a certificate in the master database				
a master key in DB1				
a database encryption key in DB1				

Answer:

Objects		Answer Area		
a database encryption key in the master database	> <	1	a master key in the master database	↑ ↓
a certificate in DB1		2	a certificate in the master database	
a master key in DB1		3	a database encryption key in DB1	

NEW QUESTION: 14

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 IO percentage. The solution must minimize costs.

What should you do?

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

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

NEW QUESTION: 15

You have 50 Azure SQL databases.

You need to automate Transact-SQL queries to run periodically against each database. The solution must minimize administrative effort.

What should you configure?

- A. elastic jobs
- B. SQL Server Agent jobs
- C. Power Automate
- D. Task Scheduler in Windows

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

The SQL Server Agent is a Microsoft Windows service that executes scheduled administrative tasks (called jobs) in SQL Server and Azure SQL Managed Instance.

Reference:

<https://learn.microsoft.com/en-us/ssms/agent/sql-server-agent>

NEW QUESTION: 16

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

You need to migrate the databases to an Azure SQL managed instance.

The solution must minimize downtime and prevent data loss.

What should you use?

A. log shipping

B. Always On availability groups

C. Database Migration Assistant

D. Backup and Restore

Answer: D (LEAVE A REPLY)

DMA does not support database migrations to Azure SQL Managed Instance. Recommendation is to use the Azure SQL migration extension for Azure Data Studio, which supports both online and offline database migrations to Azure SQL Managed Instance, but here we don't have the option "Azure SQL migration extension for Azure Data Studio".

<https://learn.microsoft.com/en-us/sql/dma/dma-overview?view=sql-server-ver16#capabilities> Regarding log shipping to Managed Instance is not possible; it only supports the restore of full backups.

<https://dba.stackexchange.com/questions/232332/is-it-possible-to-log-ship-from-on-premise-sql-server-to-azure-sql-managed-instance>

<https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-guide>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

Hotspot Question

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

1 2 4 6 8 10 12 14 16 18 20 24 32 40 6 vCores

Min vCores

0.75 1 1.25 1.5 1.75 2 2.25 2.5 3 4 5 6 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

Days 0 Hours 4 Minutes 0

Data max size 800 GB

1 GB 800 GB 1.5 TB 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.

Answer Area



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

no extra time
up to 10 minutes
up to one minute

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

intermittent and unpredictable
regular and high
steady and low

Answer:

Answer Area

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.

no extra time
up to 10 minutes
up to one minute

intermittent and unpredictable
regular and high
steady and low

The Microsoft logo, consisting of four colored squares (red, green, blue, yellow) arranged in a 2x2 grid, followed by the word "Microsoft" in a sans-serif font.

Explanation:

Box 1: no extra time

Auto Pause is not checked in the exhibit.

Note: If Auto Pause is checked the correct answer is: up to one minute

Box 2: intermittent and unpredictable

Reference:

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

NEW QUESTION: 18

You have an Azure subscription.

You plan to provision a single Azure SQL database.

You need to ensure that the database supports the autoscaling of compute resources.

Which service tier should you choose?

- A. Premium
- B. General Purpose
- C. Business Critical
- D. Standard

Answer: B (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale?view=azuresql#compare-resource-limits>

NEW QUESTION: 19

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

The output will be sent to a Delta Lake table.

Which output mode should you use?

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

Answer: B (LEAVE A REPLY)

Append Mode: Only new rows appended in the result table since the last trigger are written to external storage. This is applicable only for the queries where existing rows in the Result Table are not expected to change.

Complete Mode: The entire updated result table is written to external storage. It is up to the storage connector to decide how to handle the writing of the entire table.

Update Mode: Only the rows that were updated in the result table since the last trigger are written to external storage. This is different from Complete Mode in that Update Mode outputs only the rows that have changed since the last trigger. If the query doesn't contain aggregations, it is equivalent to Append mode.

Reference:

<https://docs.databricks.com/getting-started/spark/streaming.html>

NEW QUESTION: 20

Hotspot Question

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.

Answer Area

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:



Explanation:

Box 1: SET QUERY_STORE = ON (OPERATION MODE = READ_WRITE);

As a pre-requisite, Query Store must be turned on.

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/sql/relational-databases/automatic-tuning/automatic-tuning>

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

NEW QUESTION: 21

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.



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



What should you implement to meet the disaster recovery requirements for the PaaS solution?
A. Availability Zones
B. failover groups

- C. Always On availability groups
- D. geo-replication

Answer: B (LEAVE A REPLY)

Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.

The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.

Reference:

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

NEW QUESTION: 22

You have an Azure SQL managed instance named SQL1.

You need to configure a backup solution for SQL1. The solution must ensure that database backups can be retained for up to 10 years.

Which two resources can you include in the solution? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. a Recovery Services vault
- B. an Azure Storage account
- C. a long-term backup retention policy
- D. a policy in Azure Backup

Answer: B,C (LEAVE A REPLY)

In Azure SQL Managed Instance, you can configure a long-term backup retention policy (LTR).

This allows you to automatically retain database backups in separate Azure Blob storage containers for up to 10 years. You can then recover a database using these backups with the Azure portal, the Azure CLI and PowerShell.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/long-term-backup-retention-configure>

NEW QUESTION: 23

You have a Microsoft SQL Server 2022 instance.

You need to troubleshoot failing SQL Agent Jobs.

Which command should you run?

- A. EXEC sp_trace_setfilter
- B. EXEC sp_help_jobserver
- C. EXEC sp_help_jobsteplog
- D. EXEC sp_read_error_log;

Answer: D (LEAVE A REPLY)

sp_readerrorlog (Transact-SQL) allows you to read the contents of the SQL Server or SQL Server Agent error log file and filter on keywords.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-readerrorlog-transact-sql?>

NEW QUESTION: 24

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

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

Answer: (SHOW ANSWER)

You set up a dynamic data masking policy in the Azure portal by selecting the `Dynamic Data Masking` blade under `Security` in your `SQL Database` configuration pane. This feature cannot be set using portal for `SQL Managed Instance`. For more information, see `Dynamic Data Masking`.
<https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-overview>

NEW QUESTION: 25

You have two on premises Microsoft SQL Server 2022 instances as shown in the following table.

Name	Description
Server1	Contains a live database named DB1 that uses the simple recovery model
Server2	Used as a failover target for DB1 and contains a database named DB2

You need to configure log shipping between DB1 and DB2.
What should you do first on DB1?

- A. Create a new log backup.
- B. Run the `sp_add_log_shipping_primary_database` command.
- C. Enable the full or bulk logged recovery model.
- D. Disable SQL Server Browser.

Answer: C (LEAVE A REPLY)

Configure Log Shipping (SQL Server)
Prerequisites

*-> The primary database must use the full or bulk-logged recovery model; switching the database to simple recovery will cause log shipping to stop functioning.
* Etc.

Reference:
<https://learn.microsoft.com/en-us/sql/database-engine/log-shipping/configure-log-shipping-sql-server>

NEW QUESTION: 26

Hotspot Question
You have an Azure subscription.
You deploy resources by using the following Bicep template.

```

resource sqlServer 'Microsoft.Sql/server@2022-05-01-preview' = {
  name: 'Server1'
  location: 'Sweden central'
  properties: {
    administrators: {
      login:sqladmins
      principalType: 'Group'
      administratorType: 'ActiveDirectory'
    }
  }
}

resource sqlDB 'Microsoft.Sql/servers/databases@2022-05-01-preview' = {
  parent: 'Server1'
  name: 'Database1'
  location: 'Sweden central'
  sku: {
    name: 'Standard'
    tier: 'Standard'
  }
}

```

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

Statements	Yes	No
Server1 will support mixed authentication.	☐	☐
Database1 will scale dynamically based on the load.	☐	☐
Server1 will be created as an Azure SQL managed instance.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Server1 will support mixed authentication.	☒	☐
Database1 will scale dynamically based on the load.	☐	☒
Server1 will be created as an Azure SQL managed instance.	☐	☒

Explanation:

Box 1: Yes

The azureADOnlyAuthentication (bool) property is not present.

If set to true then Azure Active Directory only Authentication enabled.

Mixed authentication will be supported.

Box 2: No

We see sku tier is Standard.

Azure SQL databases with the Standard tier do not support autoscale. Autoscale is a feature available in the Serverless compute tier of Azure SQL Database, where resources scale automatically based on workload demands. The Standard tier, part of the provisioned compute model, requires manual scaling or the use of Elastic Pools for resource sharing and scaling.

Box 3: No

The first line of the BICEP template would have to be like:

resource managedInstance 'Microsoft.Sql/managedInstances@2021-11-01-preview' = { Reference:

<https://learn.microsoft.com/en-us/azure/templates/microsoft.sql/managedinstances?pivots=deployment-language-bicep>

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/create-bicep-quickstart>

NEW QUESTION: 27

Hotspot Question

You have a SQL Server on Azure Virtual Machines instance named SQL1.

You perform the following actions:

- Create a SQL Server Agent job named Job1.
- Configure a schedule named Schedule1.

You need to ensure that Job1 runs daily on SQL1 by using Schedule1.

Which stored procedures should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

To ensure that Job1 runs by using Schedule1:

sp_add_jobschedule

sp_attach_schedule

sp_update_schedule


To ensure that Job1 runs on SQL1:

sp_add_job

sp_add_jobserver

sp_update_job

Answer:



NEW QUESTION: 28

SIMULATION

You need to configure a monthly backup of db2. The monthly backups must be retained for five years, and then deleted.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Change automated backup settings for Azure SQL Database

You can modify automated backup settings for Azure SQL Database, such as the short-term retention policy and the backup storage redundancy option that's used for backups.

To change differential backup frequency, select 12 Hours or 24 hours from the drop-down menu under Differential backup frequency.

Step 1: Go to the logical server in Azure with the databases whose retention period you want to change.

Step 2: Select Backups on the left pane, and then select the Retention policies tab.

Step 3: Select the databases for which you want to change the PITR backup retention. [Select db2] Step 4: Select Configure policies from the action bar.

Step 5: To change the retention period for point-in-time restore backups, use the slider under Point-in-time restore. [Skip] Step 6: To change the PITR (point-in-time recovery) backup retention period or the differential backup frequency for active databases by using the Azure portal: [Skip]

Home > contosohotels-acc-sqlserver

contosohotels-acc-sqlserver | Backups

SQL server | Directory: contosohotels.com

Search

Refresh Feedback **Configure policies** Remove LTR settings

Available backups **Retention policies**

Configure and manage your automated backup retention policies below. Point-in-time enable you to keep full backups for up to 10 years. [Learn more](#)

Search for a database

Database ↑↓	PITR ↑↓	Diff. backup freq... ↑↓
Test_Database_1	7 Days	24 Hours

Configure policies

SQL server

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? ⓘ

7

Differential backup frequency
Specify how often you want differential backups to be taken. [Learn more](#)

Take a differential backup every:

24 Hours

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:

0 Week(s)

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

0 Week(s)

Apply Cancel

Step 7: In Long-term retention, Monthly LTR Backups, Set Keep the first backup of each month for: 5 Years Question: The monthly backups must be retained for five years, and then deleted.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/automated-backups-change-settings>

NEW QUESTION: 29

You have an on-premises Microsoft SQL Server 2019 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. merge
- B. snapshot
- C. transactional
- D. peer-to-peer

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

NEW QUESTION: 30

SIMULATION

You plan to perform performance testing of db1.

You need prevent db1 from reverting to the last known good query plan.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Automatic tuning options

The automatic tuning options available in Azure SQL Database and Azure SQL Managed Instance are:

- * FORCE LAST GOOD PLAN (automatic plan correction)

Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and forces queries to use the last known good plan instead of the regressed plan.

- * Etc.

You can either manually apply tuning recommendations using the Azure portal, or you can let automatic tuning autonomously apply tuning recommendations for you.

Azure defaults are set to FORCE_LAST_GOOD_PLAN enabled.

We need to disable it.

With transact-SQL:

ALTER DATABASE SET options (Transact-SQL)

Parameters include:

<automatic_tuning_option> ::=

{

AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = { DEFAULT | ON | OFF })

}

Solution with T-SQL

Stored procedure execution

You can use the SQL Server Management Studio (SSMS) user interface or Transact-SQL in an SSMS query window to execute a stored procedure. Always use the latest version of SSMS.

Use SQL Server Management Studio

Step 1: In Object Explorer, connect to an instance of SQL Server or Azure SQL Database, expand that instance, and then expand Databases.

Step 2: Expand the database that you want, expand Programmability, and then expand Stored Procedures.

Step 3: Right-click the stored procedure that you want to run [here: ALTER DATABASE SET] and select Execute Stored Procedure.

Step 4: Set the parameters: AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = { OFF }) Step 5: Select OK to execute the stored procedure. If the stored procedure doesn't have any parameters, just select OK.

Reference:

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

<https://learn.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-set-options>

<https://learn.microsoft.com/en-us/sql/relational-databases/stored-procedures/execute-a-stored-procedure>

NEW QUESTION: 31

Drag and Drop Question

You have a database named db1.

The log for db1 contains the following entry.

```
Date 10/5/2021 10:57:08 AM
Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Sources spid1595

Message
The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'
```

You need to ensure that db1 can process transactions.

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

Back up the transaction log file.

Remove db1 from the availability group.

Shrink the transaction log file.

Shrink db1.

Add db1 back to the availability group.

Answer Area

1

2

3

>

<

↑

↓

Answer:

Microsoft

Answer Area

Actions

Shrink the transaction log file.

Shrink db1.

1 Remove db1 from the availability group.

2 Back up the transaction log file.

3 Add db1 back to the availability group.

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 32

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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 he 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.



Role Assignments

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



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_writer to the new schema.
- D. For each customer, grant the customer db_writer to the existing schema.

Answer: B (LEAVE A REPLY)

Scenario: 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. Members of the db_ddladmin fixed database role can run any Data Definition Language (DDL) command in a database.

Incorrect:

Not D: db_writer does not have permissions to create database objects.


Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

NEW QUESTION: 33

Hotspot Question

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.

Answer Area



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:

Answer Area



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:

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>

NEW QUESTION: 34

Hotspot Question

You have an Azure subscription that is linked to an Azure Active Directory (Azure AD) tenant named contoso.com The subscription contains an Azure SQL database named SQL1 and an Azure web app 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 Area

CREATE USER

[App1]

[Contoso\app1]

[App1@contoso.com]

FROM

login

Windows

EXTERNAL PROVIDER

Answer:

Answer Area

CREATE USER

[App1]

[Contoso\app1]

[App1@contoso.com]

FROM

login

Windows

EXTERNAL PROVIDER

Explanation:
In T-SQL, just point to App name not managed identity.
<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-sql>

NEW QUESTION: 35

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1. You need to ensure that an automatic email notification is sent once the job completes. What should you include in the solution?

- A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
- B. From SQL Server Management Studio (SSMS), run sp_set_sqlagent_properties

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

Answer: C (LEAVE A REPLY)

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:

1. In Object Explorer, expand a SQL Server instance.
2. Right-click SQL Server Agent, and then click Properties.
3. Click Alert System.
4. Select Enable Mail Profile.
5. In the Mail system list, select Database Mail.
6. In the Mail profile list, select a mail profile for Database Mail.
7. Restart SQL Server Agent.

Note: Prerequisites include:

Enable Database Mail.

Create a Database Mail account for the SQL Server Agent service account to use.

Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.

Set the profile as the default profile for the msdb database.

Reference:

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

NEW QUESTION: 36

You have an Azure Data Factory that contains 10 pipelines.

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

What should you add to each pipeline?

- A.** an annotation
- B.** a resource tag
- C.** a run group ID
- D.** a user property
- E.** a correlation ID

Answer: A (LEAVE A REPLY)

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag.

You can define tags so you can see their performance or find errors faster.

Reference:

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

NEW QUESTION: 37

Your company has an Azure SQL database named pass4testdb. This is on an Azure SQL Server called whizlabserver1000. You have to create a contained database user in the database based on an Azure AD user.

Which of the following would you execute for this requirement? (Choose 3 answers)

- A.** Connect to the database by using the Active Directory account admin account
- B.** Create a user by using the FROM EXTERNAL PROVIDER clause

- C. Connect to the database by using a server administrator account
- D. Set the Active Directory Admin for the Azure SQL Server
- E. From the Azure portal, assign the SQL DB Contributor role to the user

Answer: A,B,D (LEAVE A REPLY)

First, you need to set the Active Directory Admin for the Azure SQL Server.
The next step is to connect to the database with the Active Directory account admin account.
And then finally create the contained database user by using the FROM EXTERNAL PROVIDER clause Option C is incorrect because you need to connect to the database using the Active Directory account admin account Option E is incorrect because you don't need to assign any RBAC role to the user
<https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-configure?tabs=azure-powershell>

NEW QUESTION: 38

Drag and Drop Question

You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures.
You need to ensure that SQL Server Agent jobs running on SQL1 can notify administrators when a failure occurs.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Create a Database Mail account.

Create a profile named AzureManagedInstance_dbmail_profile.

Create a profile named application_dbmail_profile.

Enable email notifications upon failure.

Enable pager notifications upon failure.

Answer Area

➤

➡

⬅

⬅

⬆

⬇

Microsoft

Answer:

Actions

Create a profile named application_dbmail_profile.

Enable pager notifications upon failure.

Answer Area

Create a Database Mail account.

Create a profile named AzureManagedInstance_dbmail_profile.

Enable email notifications upon failure.

Explanation:

The database mail configuration wizard would start the profile creation first, however, need to complete the account creation first before the profile creation can be completed.

NEW QUESTION: 39

You have an on-premises, 20-TB Microsoft SQL Server database named DB1 that is used by a web app named App1.

You have an Azure subscription that contains a storage account named storage1.

You need to migrate DB1 to an Azure SQL database. The solution must minimize downtime for App1.

What should you use to perform the migration?

- A. Azure Data Factory
- B. SQL Data Sync
- C. transactional replication
- D. Azure Data Box

Answer: (SHOW ANSWER)

To migrate a massive 20-TB Microsoft SQL Server database to Azure with minimal downtime while utilizing your existing Azure Storage account, you should use Transactional Replication.

While Azure Database Migration Service (DMS) is Microsoft's primary migration tool, its online (minimal downtime) migration path is strictly unsupported when the target is Azure SQL Database (it only supports online migrations to Azure SQL Managed Instance or SQL Server on Azure VMs). DMS only supports offline migrations for Azure SQL Database, which would force your 20- TB database to stay offline for a multi-day data transfer.

Transactional Replication completely bypasses this limitation:

Maximizes Uptime: It continuously replicates data changes from your on-premises "Publisher" database to the Azure SQL Database "Subscriber" in near-real-time. Your web application can remain completely active and online during the entire multi-TB data sync window.

Leverages Your Storage Account: It natively uses your existing Azure Storage account (Azure Files) as the shared replication working directory to hold and deliver snapshot files to Azure.

Handles 20 TB Scale: Because Azure SQL Database allows database sizes up to 100+ TB when using the Hyperscale service tier, Transactional Replication acts as the perfect pipeline to steadily stream data into this highly scalable layer without risking timeout failures associated with a single massive block transfer.

Reference:
<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/replication-transactional-overview>

NEW QUESTION: 40

Drag and Drop Question

You have an Azure virtual machine named Server1 that contains an instance of Microsoft SQL Server 2022 named SQL1.
SQL1 contains two databases named DB1 and DB2.

You need to take a snapshot backup of DB1 and DB2. The backup must NOT disrupt the backup chain and must NOT affect other databases.

Which three Transact-SQL commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Commands

BACKUP GROUP db1, db2
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;

BACKUP SERVER
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2), MODE = COPY_ONLY);

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2);

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;

Answer Area

>

<

^

v

 Microsoft

Answer:

Commands

```
BACKUP SERVER
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

Answer Area

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2);
```

```
BACKUP GROUP db1, db2
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF
```

NEW QUESTION: 41

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 the on-premises networks shown in the following table.

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 db1 and db2 users', '131.107.11.0',

'131.107.11.255'

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

This command sets a firewall rule at the server level (master database), allowing both DB1 and DB2 to be accessed from the Branch2 subnet. This does not meet the requirement of restricting DB1 access to Branch1 only.

NEW QUESTION: 42

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

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

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1.

You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1.

You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

A. Yes

B. No

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

Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity>

NEW QUESTION: 43

You have an instance of SQL Server on Azure Virtual Machines named SQL1.

SQL1 contains an Extended Events session named session1 that captures Microsoft SQL Server events.

You need to correlate the session events with events captured by Event Tracing for Windows (ETW).

What should you do for session1?

A. Modify the Specify Session Data Storage settings.

B. Add an action.

C. Add a target.

D. Modify the Set Session Event Filters settings.

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

NEW QUESTION: 44

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>

NEW QUESTION: 45

Hotspot Question

You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQ1.
You need to use PowerShell to enable and configure automated patching for SQ1. 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

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

Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
-SQLManagementType
```

Get-AzVMSSQLServerExtension

Set-AzVMExtension

Set-AzVMSqlServerExtension

-Full

-Lightweight

-NoAgent

Answer:

Answer Area

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

Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |

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

```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType
```

- Full
- Lightweight
- NoAgent

Explanation:

Box 1: Set-AzVMSqlServerExtension, SQL in lowercase.

Box 2: Full. SQL IAAS Extension only allows patching automatically in Full mode. Lightway gives a few metrics and NoAgent is for older versions.

<https://learn.microsoft.com/en-us/powershell/module/az.compute/set-azvmsqlserverextension?view=azps-9.7.0>

NEW QUESTION: 46

Drag and Drop Question

You have two Azure subscriptions named Subscription1 and Subscription2. Subscription1 contains an Azure SQL managed instance named MI1 and an Azure Blob Storage container named container1. MI1 contains a database named DB1. Subscription2 contains a SQL managed instance named MI2 and a Blob Storage container named container2. The virtual networks in Subscription1 and the virtual networks in Subscription2 are NOT connected.

You need to move DB1 to MI2. The solution must minimize administrative effort and downtime.

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

Restore the backup of DB1 from container2 to MI2.

Create a new Azure Data Factory pipeline that uses the Managed Virtual Network Integration Runtime. Configure a Microsoft SQL Server credential and an Azure Key Vault credential.

Configure log shipping, copy the backup to container2, and then restore the backup to MI2.

Back up DB1 to container1.

Copy the BAK files from container1 to container2.

Create a new Azure Data Factory pipeline that uses the Managed Virtual Network Integration Runtime. Configure a Microsoft SQL Server credential.

Answer Area

1

2

3

Answer:

Actions

Restore the backup of DB1 from container2 to MI2.

Create a new Azure Data Factory pipeline that uses the Managed Virtual Network Integration Runtime. Configure a Microsoft SQL Server credential and an Azure Key Vault credential.

Configure log shipping, copy the backup to container2, and then restore the backup to MI2.

Back up DB1 to container1.

Copy the BAK files from container1 to container2.

Create a new Azure Data Factory pipeline that uses the Managed Virtual Network Integration Runtime. Configure a Microsoft SQL Server credential.

Answer Area

1 Back up DB1 to container1.

2 Create a new Azure Data Factory pipeline that uses the Managed Virtual Network Integration Runtime. Configure a Microsoft SQL Server credential and an Azure Key Vault credential.

3 Configure log shipping, copy the backup to container2, and then restore the backup to MI2.

Explanation:

Step 1: Back up DB1 to container1.

Step 2: Create a new Azure Data Factory pipeline that uses Managed Virtual Network Integration runtime. Configure a Microsoft SQL Server credential and an Azure Key Vault credential.

Azure Data Factory

Because the virtual networks are not peered or connected, you cannot use online migration tools like Azure Database Migration Service (DMS) easily without setting up VPN gateways. Using Azure Data Factory (ADF) to orchestrate the .BAK file copy between the two separate Blob Storage containers is the best administrative way to bridge the network gap.

Step 3: Configure log shipping, copy the backup to container2, and then restore the backup to MI2.

Incorrect:

* Copy the BAK files from container1 to container2. + Restore the backup of DB1 from container2 to MI2.

Azure SQL Managed Instance does not support restoring a database from a .BAK file if that database is already part of a live, active instance that you want to move with minimal downtime, nor does a standard restore allow you to easily sync delta changes to keep downtime low.

* Create a new Azure Data Factory pipeline that uses Managed Virtual Network Integration runtime. Configure a Microsoft SQL Server credential.

Need the Azure Key Vault credential.

Reference:

<https://learn.microsoft.com/en-us/data-migration/sql-server/managed-instance/guide>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 47

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

You view a plan summary that shows the duration in milliseconds of each execution of query

1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the DBCC FREEPROCCACHE command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

Answer: A (LEAVE A REPLY)

As per exhibit, the execution plan 1221065 has lower execution time compared to plan 1220917.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

NEW QUESTION: 48

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements

Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

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

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

NEW QUESTION: 49

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 executes mapping data flow, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Answer: ([SHOW ANSWER](#))

Correct solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Reference:

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

NEW QUESTION: 50

Drag and Drop Question

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

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

You need to back up db1 to mysqlbackups, and then restore the backup to a new database named db2 that is hosted on SQL1. The solution must ensure that db1 is backed up to a stripe set.

Which three Transact-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

Answer Area



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

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

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

```
CREATE CREDENTIAL  
[https://mysqlbackups.blob.core.windows.net]
```



```
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

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

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

Answer:

Statements

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

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

Answer Area

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

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

```

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

```
/backups/db1_4.bak'
```

```
GO
```

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



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

Explanation:

Box 1: CREATE CREDENTIAL

[https://<

Create a credential.

The following examples create SQL Server credentials for authentication to the Microsoft Azure Blob Storage service.

CREATE CREDENTIAL

[https://<mystorageaccountname>.blob.core.windows.net/<mystorageaccountcontainername>] WITH IDENTITY = 'SHARED ACCESS SIGNATURE', SECRET = '<SAS_TOKEN>'; Incorrect:

The other CREATE CREDENTIAL option does not refer to mysqlbackups - the name of the account.

Box 2: BACKUP..

Do no include WITH CREDENTIAL when backing up To URL using Shared Access Signature.

Example:

BACKUP DATABASE AdventureWorks2016

TO URL =

'https://<mystorageaccountname>.blob.core.windows.net/<mycontainername>/AdventureWorks2016.bak';

Box 3: RESTORE

Do no include WITH CREDENTIAL when restoring up To URL using Shared Access Signature.

Example: BACKUP DATABASE AdventureWorks2016

TO URL =

'https://<mystorageaccountname>.blob.core.windows.net/<mycontainername>/AdventureWorks2016.bak';

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-server-ver15>

NEW QUESTION: 51

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource.

Which Databricks services should you log?

- A. clusters
- B. jobs
- C. DBFS
- D. SSH

E. workspace

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

Clusters logs include information regarding changes in compute.

Incorrect:

Not E: Workspace logs do not include information related to changes in compute.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/administration-guide/account-settings/azure-diagnostic-logs#configure-diagnostic-log-delivery>

NEW QUESTION: 52

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

<https://learn.microsoft.com/en-us/sql/ssms/agent/set-service-startup-account-sql-server-agent-sql-server-configuration-manager?view=sql-server-ver16>

NEW QUESTION: 53

You have an Azure SQL database named sqlldb1.

You need to minimize the amount of space by the data and log files of sqlldb1.

What should you run?

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

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

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

Incorrect Answers:

D: To shrink one data or log file at a time for a specific database, execute the DBCC SHRINKFILE command.

Reference:

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

NEW QUESTION: 54

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

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

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1. You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

A. Yes

B. No

Answer: A (LEAVE A REPLY)

This procedure is necessary and correct, as there is no option in Azure SQL Database to directly restore a database.

Reference:

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

NEW QUESTION: 55

Drag and Drop Question

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

Resize VM1.

Set Enable Ultra disk compatibility to Yes.

Stop and deallocate VM1.

Attach the ultra disk.

Start VM1.

Answer Area

Answer:

Actions



Answer Area

Stop and deallocate VM1.

Resize VM1.

Set Enable Ultra disk compatibility to **Yes**.

Attach the ultra disk.

Start VM1.

NEW QUESTION: 56

You have an Azure Database for MySQL database. You have to identify which database queries are consuming the most resources.

Which of the following tool can be used for this requirement?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

Answer: A (LEAVE A REPLY)

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns. Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance.

Reference:

<https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

NEW QUESTION: 57

Hotspot Question

You have an on-premises Microsoft SQL Server instance named SQLSVR1 that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to perform the following actions:

- Sync DB1 with SQLDB1 by using SQL Data Sync.
- Configure an index in DB1 by using CREATE INDEX.
- Add data to DB1 at regular intervals by using BULK INSERT.

Which option should you specify when you use BULK INSERT, and which parameter should you specify when you use CREATE INDEX? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Option:

- CHECK_CONSTRAINTS
- FIRE_TRIGGERS
- ROWS_PER_BATCH

Parameter:

- IGNORE_DUP_KEY
- UNIQUE
- WHERE



Answer:

Answer Area

Option:

- CHECK_CONSTRAINTS
- FIRE_TRIGGERS
- ROWS_PER_BATCH

Parameter:

- IGNORE_DUP_KEY
- UNIQUE
- WHERE



Explanation:

Using ROWS_PER_BATCH with BULK INSERT helps optimize the data loading process by indicating how many rows constitute a batch. By specifying ROWS_PER_BATCH, you can improve the performance of bulk data loading operations, which is especially beneficial when adding large amounts of data at regular intervals. It can reduce the amount of logging and improve the speed of the operation by allowing SQL Server to handle data in manageable chunks.

Specifying UNIQUE with CREATE INDEX ensures that the index enforces uniqueness on the specified column(s), which is essential if you want to prevent duplicate entries in those columns.

By ensuring uniqueness, this parameter helps maintain the quality and integrity of the data in DB1, which is important for reliable database operations and queries that depend on the indexed data.

NEW QUESTION: 58

Hotspot Question

You have an Azure SQL database named DB1.

You need to identify how much unused space in megabytes was allocated to DB1.

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

NOTE: Each correct selection is worth one point.

Answer Area



```
SELECT SUM(size/16.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB

FROM sys.database_files
sys.resource_stats
sys.dm_db_resource_stats

GROUP BY type_desc

HAVING type_desc = 'ROWS';
```

Answer:

Answer Area

```
SELECT SUM(size/16.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB

FROM sys.database_files
sys.resource_stats
sys.dm_db_resource_stats

GROUP BY type_desc

HAVING type_desc = 'ROWS';
```

Explanation:

Box 1: 128.0

-- Connect to database

-- Database data space allocated in MB and database data space allocated unused in MB SELECT SUM(size/128.0) AS DatabaseDataSpaceAllocatedInMB, SUM(size/128.0 - CAST(FILEPROPERTY(name,

'SpaceUsed') AS int)/128.0) AS DatabaseDataSpaceAllocatedUnusedInMB FROM sys.database_files GROUP BY type_desc HAVING type_desc = 'ROWS' Box 2: sys.database_files Reference:
<https://techcommunity.microsoft.com/t5/azure-database-support-blog/how-to-get-azure-sql-database-size/ba-p/369189>

NEW QUESTION: 59

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails.

The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Premium
- B. Azure SQL Database Hyperscale
- C. Azure SQL Database Basic
- D. Azure SQL Database serverless

Answer: A (LEAVE A REPLY)

To ensure automatic failover without data loss if a datacenter fails, you should choose the Business Critical deployment option and the Premium pricing tier for the Azure SQL database named DB1. The Business Critical deployment option provides a 99.995% availability SLA, and the Premium pricing tier includes features such as advanced security, high availability, and business continuity. This option also offers the ability to use a failover group with multiple readable secondary replicas.

NEW QUESTION: 60

You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.

You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.

You need to recommend a database service for the deployment. The solution must minimize administrative effort.

What should you include in the recommendation?

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

Answer: B (LEAVE A REPLY)

Azure SQL Database single database supports Data Sync.

Incorrect Answers:

A: Azure SQL Managed instance does not support Data Sync.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

NEW QUESTION: 61

Your company has an on-premises instance of Microsoft SQL Server. This instance contains a database that has a 5 TiB data file. The database requires 10,000 IOPS.

You have to port the database onto Azure. You create a virtual machine. You then allocate two 1 TiB, P30 Premium storage disks to the virtual machine. Each disk provides 5,000 IOPS. You have to ensure the storage is configured to support the migrated database.

Which of the following would you implement for this requirement? (Choose 3 answers)

- A. Create a virtual disk that uses the stripe layout
- B. Create a virtual disk that uses mirror layout

- C. Create a volume
- D. Create a storage pool

Answer: (SHOW ANSWER)

First, you need to create a storage pool.

This can be used to host multiple disks together. Now create a virtual disk based on the stripe layout.

Then create a volume from the virtual disk.

Option B is incorrect because it is normally used to mirror data from one disk onto another.

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 62

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	Azure 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 Databricks	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 Storage	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft	Yes	No	Yes	Yes

Incorrect Answers:

C: Change tracking is a lightweight solution that provides an efficient change tracking mechanism for applications. It applies to both Azure SQL Database and SQL Server.

D: Azure SQL Database and Azure SQL Managed Instance automatic tuning provides peak performance and stable workloads through continuous performance tuning based on AI and machine learning.

Reference:

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

NEW QUESTION: 63

Hotspot Question

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

 Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning.

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.

Answer Area

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:

Answer Area

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:

Box 1: No

By default CREATE INDEX is disabled. It is here configured as INHERIT so it is disabled.

Box 2: No

By default DROP INDEX is disabled.

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.

Reference:

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

NEW QUESTION: 64

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.

Hotspot Question

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Collect metrics from:

The database only

The elastic pool and the database

The elastic pool only

The server, the elastic pool, and the database

Stream metrics to:

Azure Event Hubs

Azure Log Analytics

Azure Storage

Answer:

Answer Area

Collect metrics from:

The database only

The elastic pool and the database

The elastic pool only

The server, the elastic pool, and the database

Stream metrics to:

Azure Event Hubs

Azure Log Analytics

Azure Storage

Explanation:

Box 1: The server, the elastic pool, and the database

Senario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Event hubs are able to handle custom metrics.

Incorrect Answers:

Azure Log Analytics

Azure metric and log data are sent to Azure Monitor Logs, previously known as Azure Log Analytics, directly by Azure. Azure SQL Analytics is a cloud only monitoring solution supporting streaming of

However, because Azure SQL Analytics does not use agents to connect to Azure Monitor, it does not support monitoring of SQL Server hosted on-premises or in virtual machines.

NEW QUESTION: 65

Hotspot Question

You have an Azure SQL database.

You have a query and the associated execution plan as shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

ANSWER AREA

Answer Area

**Microsoft**

The performance issue stems from the **[answer choice]** operator.

Select

Index Seek

Key Lookup

Nested Loops

The performance issue can be resolved by adding include columns to the **[answer choice]**

heap

clustered index

nonclustered index

Answer:

Answer Area

The performance issue stems from the [answer choice] operator.

Select

Index Seek

Key Lookup

Nested Loops

The performance issue can be resolved by adding include columns to the [answer choice]

heap

clustered index

nonclustered index

Explanation:

Box 1: Key Lookup

The Key Lookup cost is 99% so that is the performance bottleneck.

Box 2: nonclustered index

The key lookup on the clustered index is used because the nonclustered index does not include the required columns to resolve the query. If you add the required columns to the nonclustered index, the key lookup will not be required.

NEW QUESTION: 66

Drag and Drop Question

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 contained database user.

Answer Area

Create a resource pool.

Create a workload group.

Create a user-defined classifier function.

Modify Resource Governor.



Explanation:

1. The resource pool provides and limits the resources required by the application
2. The workload group uses the resource pool it is associated with
3. The classifier function assign the incoming session to the workload group
4. Modify resource governor to apply configuration changes

<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>

NEW QUESTION: 67

You have a Microsoft Azure virtual machine named SQL1 that has Microsoft SQL Server 2016 installed.

SQL1 has a critical database that contains online transaction data. The database is 700 GB.

You regularly perform full database backups.

On a separate Azure virtual machine, you run a test restore procedure of the database and discover that restoring the database takes longer than expected.

You need to reduce the amount of time required to restore the database.

What should you do?

- A. Create multiple filegroups in the database, and then modify the backup solution to use filegroup backups.
- B. Partition the tables in the database, and then modify the backup solution to use incremental backups.
- C. Modify the backup solution to use transaction log backups.
- D. Modify the backup solution to use differential database backups.

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

NEW QUESTION: 68

Hotspot Question

You have an Azure SQL managed instance named Server1 and an Azure Blob Storage account named storage1 that contains Microsoft SQL Server database backup files.

You plan to use Log Replay Service to migrate the backup files from storage1 to Server1. The solution must use the highest level of security when connecting to storage1.

Which PowerShell cmdlet should you run, and which parameter should you specify to secure the connection? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Cmdlet:
Complete-AzSqlInstanceDatabaseLogReplay
Get-AzSqlInstanceDatabaseLogReplay
Start-AzSqlInstanceDatabaseLogReplay

Parameter:
ManagedIdentity
SharedAccessSignature
StorageAccountKey

Microsoft

Answer:

Answer Area

Cmdlet:
Complete-AzSqlInstanceDatabaseLogReplay
Get-AzSqlInstanceDatabaseLogReplay
Start-AzSqlInstanceDatabaseLogReplay

Parameter:
ManagedIdentity
SharedAccessSignature
StorageAccountKey

Microsoft

NEW QUESTION: 69

Drag and Drop Question

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

Microsoft

Navigation icons: > < ^ v

Answer:

Actions

- Connect to DB1 by using the server administrator account.
- 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.

Microsoft

Navigation icons: > < ^ v

Explanation:

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: 70

Hotspot Question

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

You need to ensure that User1 can perform the following tasks on DB1:

- Create jobs.
- View all jobs.
- Modify, delete, and disable the jobs the user created.

The solution must use the principle of least privilege.

Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.

AnswerArea

Microsoft

Location:

DB1

master

msdb


Built-in role:

SQLAgentOperatorRole

SQLAgentReaderRole

SQLAgentUserRole

Answer:



NEW QUESTION: 71

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: B,C (LEAVE A REPLY)

On Azure SQL databases, you do not have to create a Login for the user unless he needs to access several databases within the server. The requirement is each user will only be allowed to access his own database. So with the principle of least privilege, we just need to a contained user. The database firewall is additional security to limit access on each database.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/secure-database-tutorial>

NEW QUESTION: 72

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 the on-premises networks shown in the following table.

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 DB1 and run the following command.

```
EXECUTE sp_set_firewall_rule 'Allow db1 users', '131.107.10.0',  
'131.107.10.255'
```

You connect to DB2 and run the following command.

```
EXECUTE sp_set_database_firewall_rule 'Allow db2 users',  
'131.107.11.0', '131.107.11.255'
```

Does this meet the goal?

A. Yes

B. No

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

NEW QUESTION: 73

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 shrink the transaction log file.

Does this meet the goal?

A. Yes

B. No

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

You should instead reduce the use of table variables and temporary tables.

Or you could create additional tempdb files.

Note: The following operations use tempdb extensively:

- * Repetitive create-and-drop operation of temporary tables (local or global).
- * Table variables that use tempdb for storage.
- * Etc.

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION: 74

Hotspot Question

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

CALLER
LOGIN
OWNER
USER

 = 'user1@contoso.com'

GO

SELECT SUSER_SNAME ()

REVERT
REVOKE
ROLLBACK

GO

Answer:

```
EXECUTE AS [ ] = 'user1@contoso.com'
```

CALLER
LOGIN
OWNER
USER

GO

```
SELECT SUSER_SNAME ()
```

REVERT
REVOKE
ROLLBACK



GO
Explanation:

Box 1: USER

Using EXECUTE AS and REVERT to switch context.

The following example creates a context execution stack using multiple principals. The REVERT statement is then used to reset the execution context to the previous caller.

**

```
EXECUTE AS USER = 'user2';
```

--The following REVERT statements will reset the execution context to the previous context.

```
REVERT;
```

**

Box 2: REVERT

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15>

<https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15>

NEW QUESTION: 75

Drag and Drop Question

You have SQL Server on an Azure virtual machine.

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

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

Actions

Export a built-in policy.

Create a custom policy based on a condition.

Create a custom condition based on a built-in facet.

View the policy history.

Import a policy file.

Run a policy evaluation.

Answer Area

⬅

➡

⬆

⬇

Answer:

Actions

Export a built-in policy.

View the policy history.

Import a policy file.

Answer Area

Create a custom condition based on a built-in facet.

Create a custom policy based on a condition.

Run a policy evaluation.

⬅

➡

⬆

⬇

Explanation:

First create a condition, then a custom policy based on the condition, finally run a policy evaluation.

Reference:

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

NEW QUESTION: 76

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. an Azure Standard Load Balancer
- B. a virtual network name (VNN) resource
- C. a Basic Azure Load Balancer

D. a distributed network name (DNN) resource

Answer: D (LEAVE A REPLY)

You should create a distributed network name (DNN) resource 1. On Azure Virtual Machines, the DNN routes traffic to the appropriate clustered resource. It provides an easier way to connect to the SQL Server failover cluster instance (FCI) than the virtual network name (VNN), without the need for an Azure Load Balancer 1.

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 77

SIMULATION

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.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

View or change the compatibility level of a database

You can view or change the compatibility level of a database in SQL Server, Azure SQL Database, or Azure SQL Managed Instance by using SQL Server Management Studio or Transact-SQL.

Use SQL Server Management Studio

To view or change the compatibility level of a database using SQL Server Management Studio (SSMS) Step 1: Connect to the appropriate server or instance hosting your database [Here db1].

Step 2: Select the server name in Object Explorer.

Step 3: Expand Databases, and, depending on the database, either select a user database or expand System Databases and select a system database.

Step 4: Right-click the database, and then select Properties.

The Database Properties dialog box opens.

Step 5: In the Select a page pane, select Options.

The current compatibility level is displayed in the Compatibility level list box.

Step 6: To change the compatibility level, select a different option from the list [Select SQL Server 2012]

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/databases/view-or-change-the-compatibility-level-of-a-database>

NEW QUESTION: 78

Drag and Drop Question

You are building an Azure virtual machine.

You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.

You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.

You need to configure storage for the virtual machine to support the database.

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

Actions

a virtual disk that uses the stripe layout

a virtual disk that uses the mirror layout

a volume

a virtual disk that uses the simple layout

a storage pool

Answer Area



Answer:

Actions

a virtual disk that uses the stripe layout

a virtual disk that uses the mirror layout

Answer Area

a storage pool

a virtual disk that uses the simple layout

a volume



Explanation:

Follow these same steps to create striped virtual disk:

- Create Log Storage Pool.
- Create Virtual Disk
- Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses the simple layout

There are only three layouts while creating virtual disks on a storage pool:

Simple - Data is striped across physical disks, maximizing capacity and throughput Mirror - Data is striped across physical disks, creating two or three copies of data Parity - Data and parity information is striped across physical disks, increasing reliability Box 3: a volume Reference:

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

NEW QUESTION: 79

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A.** distributed availability groups
- B.** database mirroring
- C.** log shipping
- D.** Azure Database Migration Service

Answer: D (LEAVE A REPLY)

When you migrate databases to Azure by using Azure Database Migration Service, you can do an offline or an online migration. With an offline migration, application downtime starts when the migration starts.

With an online migration, downtime is limited to the time to cut over at the end of migration. We suggest that you test an offline migration to determine whether the downtime is acceptable; if not, do an online migration.

Reference:

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

NEW QUESTION: 80

Drag and Drop Question

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1.

You need to encrypt the data in Col1.

Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

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

Answer Area



Answer:

Actions

- Create a column master key.
- Update Col1.

Answer Area

- Create a database master key.
- Create a certificate.
- Create a symmetric key.
- Open the symmetric key.

Explanation:

Use the following steps for column level encryption:

1. Create a database master key (Step 1)

- 2. Create a self-signed certificate for SQL Server (Step 2)
- 3. Configure a symmetric key for encryption (Step 3)
- 4. Encrypt the column data (this includes Open the symmetric key - Step 4)
- 5. Query and verify the encryption

Step 1: Create a database master key

Create a database master key for column level SQL Server encryption

In this first step, we define a database master key and provide a password to protect it. It is a symmetric key for protecting the private keys and asymmetric keys.

Step 2: Create a certificate.

Create a self-signed certificate for Column level SQL Server encryption In this step, we create a self-signed certificate using the CREATE CERTIFICATE statement. You might have seen that an organization receives a certificate from a certification authority and incorporates into their infrastructures. In SQL Server, we can use a self-signed certificate without using a certification authority certificate.

Step 3: Create a symmetric key.

Configure a symmetric key for column level SQL Server encryption.

In this step, we will define a symmetric key that you can see in the encryption hierarchy as well.

The symmetric key uses a single key for encryption and decryption as well.

Step 4: Open the symmetric key

Data encryption.

Let's encrypt the data in this newly added column.

In a query window, open the symmetric key and decrypt using the certificate. We need to use the same symmetric key and certificate name that we created earlier Etc.

Reference:

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

NEW QUESTION: 81

SIMULATION

You need to implement a disaster recovery solution by using active geo-replication for an Azure Azure SQL database named db1. The replica must be in the East US Azure region on a server named sql23456789-dr.database.windows.net.

You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Step 1: Navigate to your database: In the Azure portal, browse to the Azure SQL Database you want to replicate. [Select db1] Step 2: On the SQL Database page, select your database, scroll to Data management, select Replicas, and then select Create replica.



Step 3: Select your geo-secondary database Subscription and Resource group.

Geo-secondary database details

Project details

Select the subscription of geo secondary database to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ

DS-SQLDB

Resource group * ⓘ

AzureSQLDB

Create new

Step 4: Select or create the server for the secondary database [Select named sql23456789- dr.database.windows.net] Question: The replica must be in the East US Azure region on a server named sql23456789- dr.database.windows.net.

Step 5: When secondary database is created [or selected], the deployment process begins.

DeleteCancelRedeployRefresh

Deployment is in progress

Deployment name: Microsoft.SQLDatabase.newDatabaseCopyExisti...
Subscription: Contoso
Resource group: ContosoHotel

Start time: 8/20/2021, 3:16:00 PM
Correlation ID:

Deployment details (Download)

Resource	Type	Status	Operation details
contosowest/guestlist	Microsoft.Sql/servers/databases	Accepted	Operation details

Step 6: When the deployment is complete, the secondary database displays its status.

✓ Your deployment is complete

 Deployment name: Microsoft.SQLDatabase.newDatabaseCopyExisti... Start time: 8/20/2021, 3:16:00 PM
Subscription: Contoso Correlation ID:
Resource group: ContosoHotel

^ Deployment details [\(Download\)](#)

Resource	Type	Status	Operation details
✓ contosowest/guestlist	Microsoft.Sql/servers/databases	Created	Operation details

(Step 7: Return to the primary database page, and then select Replicas. Your secondary database is listed under Geo replicas.) Reference:
<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-configure-portal>

NEW QUESTION: 82

Hotspot Question

You have an Azure virtual machine named Server1 that has Microsoft SQL Server installed.

Server1 contains a database named DB1.

You have a logical SQL server named ASVR1 that contains an Azure SQL database named ADB1.

You plan to use SQL Data Sync to migrate DB1 from Server1 to ASVR1.

You need to prepare the environment for the migration. The solution must ensure that the connection from Server1 to ADB1 does NOT use a public endpoint.

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

NOTE: Each correct selection is worth one point.

Answer Area

Sync data by using:

Set the database isolation level to:

Answer:
Answer Area

Sync data by using:

Set the database isolation level to:

NEW QUESTION: 83

SIMULATION

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:

Solution 1 [Less complex]

Configure column encryption using Always Encrypted Wizard

The Always Encrypted Wizard is a powerful tool that allows you to set the desired Always Encrypted configuration for selected database columns. Depending on the current configuration and the desired target

configuration, the wizard can encrypt a column, decrypt it (remove encryption), or re-encrypt it (for example, using a new column encryption key or an encryption type that is different from the current type, configured for the column). Multiple columns can be configured in a single run of the wizard.

Step 1: Open the Always Encrypted Wizard

You can launch the wizard at three different levels:

At a database level - if you want to encrypt multiple columns located in different tables.

At a table level - if you want to encrypt multiple columns located in the same table.

*-> At a column level - if you want to encrypt one specific column. [This option here] Step 2: Connect to your SQL Server with the Object Explorer component of SQL Server Management Studio.

Step 3: To encrypt an individual column, navigate to the column, right-click on it, and then select Encrypt Columns. [Select the LastName column of the SalesLT.Customer table in db1] Step 4: Column

Selection page In this page, you select columns you want to encrypt, re-encrypt, or decrypt, and you define the target encryption configuration for the selected columns.

Step 5: If you want the wizard to encrypt or re-encrypt one or more columns using a new column encryption key, pick a key containing (New) in its name. The wizard will generate the key.

Step 6: Master Key Configuration page

If you have selected an autogenerated column encryption key for any column on the previous page, in this page you need to either select an existing column master key or configure a new column master key that will encrypt the column encryption key.

When configuring a new column master key, you can either pick an existing key in Windows Certificate Store or in Azure Key Vault and have the wizard to create just a metadata object for the key in the database, or you can choose to generate both the key and the metadata object describing the key in the database.

Step 7: Click "Next" and follow the wizard's instructions to complete the process.

Solution 2 [A bit more complex. Use for reference]

To encrypt a specific column using Always Encrypted with a Windows Certificate Store, you'll first need to set up a Column Master Key (CMK) [Task 1] and a Column Encryption Key (CEK) [Task 2]. The CMK will be protected by a certificate in the Windows Certificate Store, and the CEK will then be used to encrypt the data in your chosen column.

Task 1: Set up a Column Master Key (CMK)

Create a self-signed certificate using SQL Server Management Studio (SSMS) Provision Column Master Keys with the New Column Master Key Dialog The New Column Master Key dialog allows you to generate a column master key or pick an existing key in a key store, and create column master key metadata for the created or selected key in the database.

Step 1: Using Object Explorer, navigate to the Security -> Always Encrypted Keys node under your database.

Step 2: Right-click on the Column Master Keys node and select New Column Master Key....

Step 3: In the New Column Master Key dialog, enter the name of the column master key metadata object.

Step 4: Select a key store:

*-> Certificate Store - Current User - indicates the Current User certificate store location in the Windows Certificate Store, which is your personal store. [Select this one] Certificate Store - Local computer Azure Key Vault Etc.

Step 5: If you're using SQL Server 2019 (15.x) and your SQL Server instance is configured with a secure enclave, you can select the Allow enclave computations checkbox to make the master key enclave-

enabled. [Skip] Step 6: Pick an existing key in your key store, or click the Generate Key or Generate Certificate button, to create a key in the key store. [Pick existing] Step 7: Click OK and the new key appears in the list.

Task 2: Create a Column Encryption Key (CEK)

1. In SSMS, navigate to "Always Encrypted Keys" -> "Column Encryption Keys".

2. Right-click and select "New Column Encryption Key".

3. Give the CEK a name.

4. Select the CMK you just created. [Created in Stage 1]

5. Click "OK".

Task 3: Encrypt the data in your chosen column.

1. Right-click on the specific column you want to encrypt in Object Explorer. [Select the LastName column of the SalesLT.Customer table in db1]

2. Select "Tasks" -> "Encrypt Columns".

3. The Always Encrypted Wizard will guide you through the process:

4. Select the column you want to encrypt. [Select the LastName column of the SalesLT.Customer table in db1]
5. Choose the encryption type (Deterministic or Randomized).
6. Select the CEK you created. [Created in Stage 2]
7. Click "Next" and follow the wizard's instructions to complete the process.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted- wizard>

<https://learn.microsoft.com/en-us/sql/relational-databases/security/encryption/create-and-store- column-master-keys-always-encrypted>

NEW QUESTION: 84

Hotspot Question

You have an Azure subscription that contains a SQL Server instance on a virtual machine named Server1.

You execute the following T-SQL statement on Server1.

```
EXEC msdb.dbo.sp_add_alert @name=N'Fatal Error',
    @message_id=0,
    @severity=20,
    @enabled=1,
    @delay_between_responses=0,
    @include_event_description_in=1;
GO

EXEC msdb.dbo.sp_add_operator
    @name = 'DBA-Team',
    @enabled = 1,
    @email_address = 'DBA-Team@contoso.com';
GO

EXEC msdb.dbo.sp_add_notification
    @alert_name = N'Fatal Error',
    @operator_name = 'DBA-Team',
    @notification_method = 1;
GO
```

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

NOTE: Each correct selection is worth one point.

Answer Area



Statements

	Yes	No
An alert will be generated when a fatal error occurs.	☐	☐
A SQL Server Agent job will run when an alert is generated.	☐	☐
An email will be sent to DBA-Team@contoso.com when an alert is generated.	☐	☐

Answer:

Answer Area



Statements

	Yes	No
An alert will be generated when a fatal error occurs.	☒	☐
A SQL Server Agent job will run when an alert is generated.	☐	☒
An email will be sent to DBA-Team@contoso.com when an alert is generated.	☒	☐

NEW QUESTION: 85

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.

Hotspot Question

You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure.

The solution must meet the business requirements.
What should you include in the recommendation? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area



Quorum model:

Cloud witness

Disk witness

File share witness

Azure resource for the availability group listener:

Azure Application Gateway

Azure Basic Load Balancer

Answer:

Answer Area



Quorum model:

Cloud witness

Disk witness

File share witness

Azure resource for the availability group listener:

Azure Application Gateway

Azure Basic Load Balancer

Explanation:

Scenario: Business Requirements

Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.

Box 1: Cloud witness

If you have a Failover Cluster deployment, where all nodes can reach the internet (by extension of Azure), it is recommended that you configure a Cloud Witness as your quorum witness resource.

Box 2: Azure Basic Load Balancer

Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a public Virtual IP (VIP) that is accessible over the internet. An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.

Reference:

<https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness>

https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

You have an Azure subscription that contains multiple Azure SQL databases.
You plan to deploy automation tasks that will be controlled by using elastic jobs.
You need to ensure that the jobs authenticate by using a user-assigned managed identity.
What should you do first?

- A. Configure the minimum TLS setting on the job target logical servers for Azure SQL Database.
- B. Create a database scoped credential on the job target logical servers for Azure SQL Database.
- C. Enable Microsoft Entra authentication on the job target logical servers for Azure SQL Database.
- D. Configure an elastic jobs private endpoint.

Answer: (SHOW ANSWER)

Azure SQL Database, Create, configure, and manage elastic jobs
There are two methods of authentication for the job agent to target server(s)/database(s), Microsoft Entra authentication with a user-assigned managed identity (UMI), or database-scoped credentials.
Reference:
<https://learn.microsoft.com/en-us/azure/azure-sql/database/elastic-jobs-tutorial>

NEW QUESTION: 87

Hotspot Question
You have an Azure SQL database named SQL1.
You need to monitor the resource usage of SQL1 by using a dynamic management view. The solution must return the average memory usage and the peak CPU usage.
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

SELECT database_name = SQL1(),
AVG(avg_memory_usage_percent) AS 'Average memory usage'
MAX(

avg_cpu_percent
cpu_limit
max_worker_percent

) AS 'Maximum CPU usage'

FROM sys.

dm_db_resource_stats
dm_db_wait_stats
dm_operation_status

Answer Area



```
SELECT database_name = SQL1(),  
       AVG(avg_memory_usage_percent) AS 'Average memory usage'  
       MAX(  ) AS 'Maximum CPU usage'  
  
FROM sys. 
```

Explanation:

Box 1: max_worker_percent

Maximum concurrent workers (requests) in percentage of the limit of the database's service tier.

Box 2: dm_db_resource_stats:

This dynamic management view (DMV) in Azure SQL Database provides real-time resource utilization data.

The query `SELECT avg_memory_usage_percent FROM sys.` is used in Azure SQL Database to retrieve the average memory usage percentage of a database. This information is provided by the `sys.dm_db_resource_stats` dynamic management view (DMV). The `avg_memory_usage_percent` column represents the average memory usage as a percentage of the maximum allowed memory for the database's service tier and performance level.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-resource-stats-azure-sql-database>

NEW QUESTION: 88

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.

What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the DBCC DBREINDEX command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Answer: (SHOW ANSWER)

Automatic tuning is a fully managed intelligent performance service that uses built-in intelligence to continuously monitor queries executed on a database, and it automatically improves their performance.

Automatic tuning for Azure SQL Database uses the CREATE INDEX, DROP INDEX, and FORCE LAST GOOD PLAN database advisor recommendations to optimize your database performance.

Reference:

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

NEW QUESTION: 89

Hotspot Question

You have an Azure subscription that contains an Azure SQL database named DB1.

You execute the following T-SQL statements against DB1.

```
CREATE TABLE table1 (  
    column1 TINYINT,  
    column2 VARCHAR(255),  
    column3 NCHAR(10),  
    column4 BINARY(50),  
    column5 DATETIME2  
);
```

```
ALTER TABLE table1 REBUILD WITH (DATA COMPRESSION = ROW)
```

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

NOTE: Each correct selection is worth one point.

Answer Area		
Statements	Yes	No
All the fields in column3 will occupy 10 bytes on disk.	☐	☐
Dictionary compression will be used to reduce the size of table1 on disk.	☐	☐
The data types stored in column1, column2, and column4 support compression.	☐	☐

Answer:

Statements

All the fields in column3 will occupy 10 bytes on disk.

Yes

☒

No

☐

Dictionary compression will be used to reduce the size of table1 on disk.

☐
☒

The data types stored in column1, column2, and column4 support compression.

☒
☐

Explanation:

The NCHAR(10) data type is a fixed-length Unicode character data type. Each character in NCHAR takes 2 bytes because it uses UTF-16 encoding. Therefore, NCHAR(10) will always occupy 20 bytes on disk (2 bytes per character for 10 characters), not 10 bytes. However, the statement is asking if each field occupies 10 bytes, which is misleading; it should be 20 bytes.

The statement "ALTER TABLE table1 REBUILD WITH (DATA_COMPRESSION = ROW);" specifies row-level compression, not dictionary compression. Row compression reduces the size of the data by storing fixed-length data types in variable-length storage format and by reducing metadata overhead. Dictionary compression is a feature of page compression, not row compression.

Column1 (TINYINT): Supports compression. Row compression can store this as a smaller size if possible. Column2 (VARCHAR(255)): Supports compression. Row compression can store only the used length of the string. Column4 (BINARY(50)): Supports compression. Row compression can store binary data more efficiently by removing unnecessary padding.

NEW QUESTION: 90

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

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network.

Exposing your service to the public internet is no longer necessary.

Reference:

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

NEW QUESTION: 91

You have an Azure SQL managed instance named SQLMI1 that has the following settings:

- vCores: 4
- Service tier: General Purpose
- Hardware generation: Standard-series (Gen5)

You discover that memory pressure on SQLMI1 is high.

You need to reduce the memory pressure on SQLMI1. The solution must minimize costs.

What should to do?

- A.** Enable the Query Store.
- B.** Change vCores to 8.
- C.** Change Hardware generation to Premium-series.
- D.** Change Service tier to Business Critical.

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

The Query Store option is a cost-effective way to reduce memory pressure on a managed Azure SQL instance. The Query Store stores temporary data used by queries in the cloud, reducing the memory requirements of the managed Azure SQL instance.

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 92

Hotspot Question

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.

Answer Area

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD



Answer:

Answer Area

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD



Explanation:

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.

Incorrect Answers:

NOINIT Indicates that the backup set is appended to the specified media set, preserving existing backup sets. If a media password is defined for the media set, the password must be supplied. NOINIT is the default.

UNLOAD

Specifies that the tape is automatically rewound and unloaded when the backup is finished.

UNLOAD is the default when a session begins.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online>

NEW QUESTION: 93

You have an Azure subscription that contains a Standard SKU Azure SQL database named DB1.

You need to enable point-in-time recovery (PITR) backups for DB1. The solution must maximize how long the backups are retained.

Which retention period should you configure?

- A. seven days
- B. 35 days
- C. 180 days
- D. 365 days

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

Backup storage (point-in-time restore)

By default, seven days of database backups are stored in RA-GRS Standard blob storage. The retention period can be increased to up to 35 days.

The Azure SQL Database service protects all databases with an automated backup system.

These backups are retained for 7 days for Basic, 35 days for Standard and 35 days for Premium.

Reference:

<https://azure.microsoft.com/en-us/pricing/details/azure-sql-database/single>

<https://azure.microsoft.com/en-us/blog/azure-sql-database-point-in-time-restore/>

NEW QUESTION: 94

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>

NEW QUESTION: 95

Hotspot Question

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.

Answer Area

Five-year-old data:

Delete the blob.

Move to archive storage.

Move to cool storage.

Move to hot storage.

Seven-year-old data:

Delete the blob.

Move to archive storage.

Move to cool storage.

Move to hot storage.

Answer:

Answer Area

Five-year-old data:

Delete the blob.

Move to archive storage.

Move to cool storage.

Move to hot storage.

Seven-year-old data:

Delete the blob.

Move to archive storage.

Move to cool storage.

Move to hot storage.

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: 96

Hotspot Question

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 area.

NOTE: Each correct selection is worth one point.

Answer Area

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

@allowed()
@secure()
@sys
param administratorLoginPassword string

sqlServer 'Microsoft.Sql/servers@2022-05-01-preview' = {
  output
  param
  resource
  name: serverName
  location: location
  properties: {
    administratorLogin: administratorLogin
    administratorLoginPassword: administratorLoginPassword
  }
}
```



Answer:

Answer Area

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

```
param location string = resourceGroup().location
```

```
param administratorLogin string
```

@allowed()

@secure()

@sys

```
param administratorLoginPassword string
```

output

param

resource

```
sqlServer 'Microsoft.Sql/servers@2022-05-01-preview' = {
```

```
  name: serverName
```

```
  location: location
```

```
  properties: {
```

```
    administratorLogin: administratorLogin
```

```
    administratorLoginPassword: administratorLoginPassword
```

```
}
```



NEW QUESTION: 97

You are planning a solution that will use Azure SQL Database. Usage of the solution will peak from October 1 to January 1 each year.

During peak usage, the database will require the following:

- 24 cores
- 500 GB of storage
- 124 GB of memory
- More than 50,000 IOPS

During periods of off-peak usage, the service tier of Azure SQL Database will be set to Standard.

Which service tier should you use during peak usage?

A. Business Critical

- B. Premium
- C. Hyperscale

Answer: B (LEAVE A REPLY)

In the DTU Purchase Model with "Standard". Business Critical is with General Purpose in vCore Purchase Model.

Migrating a database from the DTU-based purchasing model to the vCore-based purchasing model is similar to scaling between service objectives in the Basic, Standard, and Premium service tiers, with similar duration and a minimal downtime at the end of the migration process. A database migrated to the vCore-based purchasing model can be migrated back to the DTU- based purchasing model at any time in the same fashion, with the exception of databases migrated to the Hyperscale service tier.

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-dtu-to-vcore?view=azuresql>

NEW QUESTION: 98

Drag and Drop Question

You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1.

You plan to deploy an Azure web app named webapp1 that will export the rows in Table1 that have changed.

You need to ensure that webapp1 can identify the changes to Table1. The solution must meet the following requirements:

- Minimize compute times.
- Minimize storage.

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

Actions

From webapp1, connect to SQLDb1, obtain the initial dataset and run the `CHANGE_TRACKING_CURREWT_VERSION()` function.

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_table`

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_db`

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`

Answer Area

Answer:

Actions

From webapp1, connect to SQLDb1, obtain the initial dataset and run the `CHANGE_TRACKING_CURRENT_VERSION()` function.

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_table`

Connect to SQLDb1 and run the following Transact-SQL statement
`EXEC sys.sp_cdc_enable_db`

Answer Area

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON`

Connect to SQLDb1 and run the following Transact-SQL statement
`ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING`

From webapp1, connect to SQLDb1, obtain the initial dataset, and run the `CHANGETABLE()` function.



Explanation:

<https://learn.microsoft.com/en-us/sql/relational-databases/track-changes/about-change-tracking-sql-server?view=sql-server-ver16> To track changes, change tracking must first be enabled for the database and then enabled for the tables that you want to track within that database. The table definition does not have to be changed in any way, and no triggers are created.

NEW QUESTION: 99

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: B (LEAVE A REPLY)

Azure Synapse Analytics does not support R script.

Correct solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/r-developers-guide>

NEW QUESTION: 100

SIMULATION

You need to classify the StandardCost and ListPrice columns of the SalesLT.Product table in db1. The columns must have the following configurations:

- An information type of Financial
- A sensitivity label of Confidential


To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:


To classify your SQL Server database:

Step 1: In SQL Server Management Studio (SSMS), connect to the SQL Server.

Step 2: In the SSMS Object Explorer, select the database [Here db1] that you would like to classify and choose Tasks > Data Discovery and Classification > Classify Data....



Step 3: The classification engine scans your database for columns containing potentially sensitive data and provides a list of recommended column classifications:



Step 4: You can change the recommended Information Type and Sensitivity Label using the drop- down boxes.

31 columns with classification recommendations (click to minimize) Microsoft

Save selected recommendations **Dismiss selected recommendations**

☐	Schema	Table	Column	Information Type	Se
☐	dbo	ErrorLog	UserName	Credentials	
☒	HumanResources	Employee	BirthDate	Date Of Birth	
☒	HumanResources	Employee	NationalIDNumber	National ID	

Step 5: Select columns StandardCost and ListPrice. Unselect other columns if necessary. For StandardCost and ListPrice columns select information type of Financial, and sensitivity label of Confidential.

Step 6: To apply the selected recommendations, select the Save selected recommendations button.

Save selected recommendations

Step 7: To display the classified columns, select appropriate schema and corresponding table from the drop-down, then select Load Columns.

Save **Add Classification** **View Report** **Refresh**

✓ The classification changes have been updated successfully.

i We have found 25 columns with classification recommendations. Click here to view them.

Choose a table to view its classified columns:

Schema **HumanResources** Table **Employee** **Load Columns**

Column	Information Type	Sensitivity Label
BirthDate	Date Of Birth	Confidential
NationalIDNumber	SSN	Confidential

Step 8: You can also manually classify columns as an alternative, or in addition, to the recommendation-based classification:

Select Add classification in the top menu of the window.

Data Classification - ShiraDB **Save** **Add Classification** **View Report**

Step 9: If you want to add classification for all the unclassified columns for a specific table in a single attempt, then select All Unclassified in the Column drop down of Add Classification page.

Add Classification Microsoft
Data Classification

Adding classification for HumanResources.Employee


Column
All Unclassified

Information Type
Other

Sensitivity Label
Confidential

Add Cancel

Step 10: To complete your classification and persistently label (tag) the database columns with the new classification metadata, select the Save button on the top menu of the window.



Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/security/sql-data-discovery-and-classification>

NEW QUESTION: 101

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: C,E (LEAVE A REPLY)

Row-Level Security enables you to use group membership or execution context to control access to rows in a database table. Implement RLS by using the CREATE SECURITY POLICY Transact-SQL

statement, and predicates created as inline table-valued functions.

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>

<https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15>

NEW QUESTION: 102

Hotspot Question

You have an Azure subscription that contains an Azure SQL database.

The database fails to respond to queries in a timely manner.

You need to identify whether the issue relates to resource_semaphore waits.

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

NOTE: Each correct selection is worth one point.

```
SELECT
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    JOIN sys.dm_exec_sessions AS dm1
    ON dm1.session_id = dm2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Answer:

```
SELECT
    is_user_process
    wait_time
    wait_type
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Explanation:

If your top wait type is RESOURCE_SEMAPHORE and you don't have a high CPU usage issue, you may have a memory grant waiting issue.

Determine if a RESOURCE_SEMAPHORE wait is a top wait

Use the following query to determine if a RESOURCE_SEMAPHORE wait is a top wait
SELECT wait_type, SUM(wait_time) AS total_wait_time_ms FROM sys.dm_exec_requests AS req JOIN sys.dm_exec_sessions AS sess ON req.session_id = sess.session_id WHERE is_user_process = 1 - GROUP BY wait_type - ORDER BY SUM(wait_time) DESC; Reference:

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

NEW QUESTION: 103

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: C (LEAVE A REPLY)

To ensure that all tables in a new Azure SQL database automatically provide cryptographic proof when data is altered and allow for auditor verification, you must enable the Ledger feature at the database level.

T-SQL to Create a Ledger Database

When creating the database, use the following T-SQL statement:

```
CREATE DATABASE [YourDatabaseName]
WITH LEDGER = ON;
```

Key Features of a Ledger Database

Automatic Enforcement: When LEDGER = ON is specified at the database level, all future tables created in that database will be ledger tables by default, even if you do not explicitly include the LEDGER = ON clause in your CREATE TABLE statements.

Tamper-Evidence: This feature uses centralized blockchain technology to cryptographically hash and link all transactions, making any unauthorized data modification evident to auditors.

Auditor Verification: Auditors can use database digests-stored in a secure, immutable location like Azure Confidential Ledger-to verify that the database's current state matches the cryptographically signed history.

Table Types: By default, new tables are created as updatable ledger tables, which maintain a history of all changes (inserts, updates, and deletes) in a transparently managed history table.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/security/ledger/ledger-overview>

NEW QUESTION: 104

You have a Microsoft SQL Server 2019 database named DB1 and an Azure SQL managed instance named SQLMI1.

You need to move a SQL Server Agent job from DB1 to SQLMI1.

Which job attribute is unsupported in SQLMI1?

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

Answer: (SHOW ANSWER)

<https://learn.microsoft.com/en-us/answers/questions/122544/azure-sql-managed-instance-sql-agent-job-output-st>

NEW QUESTION: 105

Hotspot Question

You have an Azure subscription that contains an Azure SQL managed instance named SQL1.

You use the Microsoft Power BI service.

You need to configure connectivity from the Power BI service to SQL1. The solution must ensure that only the Power BI service can initiate connections to SQL1.

Which type of endpoint should you use for SQL1, and what should you use to control the connections to SQL1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Microsoft

Endpoint type:

▼

Private

Public

Service

Control connections by using:

▼

A database-level firewall

A network security group (NSG)

A server-level firewall

Virtual network firewall rules

Answer:

Answer Area

Endpoint type:

▼

Private

Public

Service

Control connections by using:

▼

A database-level firewall

A network security group (NSG)

A server-level firewall

Virtual network firewall rules

Explanation:

To configure connectivity from the Power BI service to an Azure SQL managed instance, you should use a Public endpoint for SQL1 1. You can control the connections to SQL1 by using a Network Security Group (NSG) 1. Therefore, the correct answer is Endpoint type: Public and Control connections by using: A network security group (NSG).

NEW QUESTION: 106

You have an Azure SQL database named DB1.

You need to query the fragmentation information of data and indexes for the tables in DB1.

Which command should you run?

- A. sts.dm_db_index_physical_stats
- B. DBCC SHOWCONTIG
- C. DBCC CHECKALLOC
- D. sys.dm_db_index_usage_stats

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

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 107

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

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

You have SQL Server 2019 on an Azure virtual machine.

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

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

You need to improve system performance.

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

Does this meet the goal?

- A. Yes
- B. No

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

The following operations use tempdb extensively:

- * Repetitive create-and-drop operation of temporary tables (local or global).
- * Table variables that use tempdb for storage.
- * Etc.

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION: 108

You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1.

You have an Azure subscription that contains an

Azure SQL managed instance named MI1 and an Azure Storage account named storage1.

You plan to migrate db1 to MI1 by using the backup and restore process.

You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:

- Use block blob storage.
- Maximize security.

What should you do on storage1?

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

Answer: ([SHOW ANSWER](#))

For backup and restore process first you need to backup to Azure Storage and in order to authenticate to Azure Storage first you need to create SAS through which authentication would be done and .bak file would be created on Azure storage.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-guide?view=azuresql#backup-and-restore>

NEW QUESTION: 109

You have an application that uses a Microsoft SQL Server database on a Microsoft Azure virtual machine.

The application experiences performance issues, which you suspect are related to the connection pooling.

The issues are prevalent only when there are more than 150 concurrent connections.

You need to identify whether the performance issues are related to the connection pooling.

Which dynamic management view should you query?

- A.** sys.dm_exec_sessions
- B.** sys.dm_exec_requests
- C.** sys.dm_os_memory_pools
- D.** sys.dm_exec_connections

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

NEW QUESTION: 110

You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool.

The date dimension table will be used by all the fact tables.

Which distribution type should you recommend to minimize data movement?

- A.** ROUND_ROBIN
- B.** HASH
- C.** REPLICATE

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

NEW QUESTION: 111

Drag and Drop Question

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

A recent failure causes the master database to be lost.

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

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

NEW QUESTION: 112

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.

Correct solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Reference:

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

NEW QUESTION: 113

Hotspot Question

You have an Azure subscription.


You plan to deploy an Azure SQL database by using an Azure Resource Manager 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





```
{
  "resources": [
    {
      "type": 
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
    },
    {
      "resources": [
        {
          "type": "databases",
          "apiVersion": "2020-02-02-preview",
          ...
        },
        {
          "dependsOn": [
            "properties": [
              "tags": [
                "[resourceId('Microsoft.Sql/servers', concat(parameters('name1')))]"
              ]
            ]
          ]
        },
        ...
      ]
    }
  ]
}
```

Answer:

Answer Area



```
{
  "resources": [
    {
      "type": 
      "apiVersion": "2020-02-02-preview",
      "name": "[parameters('name1')]",
      "location": "[parameters('location')]",
      ...
    }
  ],
  ...
  "resources": [
    {
      "type": "databases",
      "apiVersion": "2020-02-02-preview",
      ...
    }
  ],
  ...
  
  "[resourceId('Microsoft.Sql/servers', concat(parameters('name1')))]"
  ...
}
```

Explanation:

Box 1: "Microsoft.Sql/servers"

Example:

```
"resources": [  
  {  
    "type": "Microsoft.Sql/servers",  
    "apiVersion": "2021-08-01-preview",  
    "name": "[parameters('serverName')]",  
    "location": "[parameters('location')]",  
    "properties": {  
      "administratorLogin": "[parameters('administratorLogin')]",  
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]"  
    },  
    {  
      "type": "Microsoft.Sql/servers/databases",  
      "apiVersion": "2021-08-01-preview",  
      "name": "[format('{0}/{1}', parameters('serverName'), parameters('sqlDBName'))]",  
      "location": "[parameters('location')]",  
      "sku": {  
        "name": "Standard",  
        "tier": "Standard"  
      },  
      "dependsOn": [  
        "[resourceId('Microsoft.Sql/servers', parameters('serverName'))]"  
      ]  
    }  
  ]  
}
```

Box 2: "dependsOn": [

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart>

NEW QUESTION: 114

Your company has an Azure SQL database named pass4testdb. Below are the configuration settings set for the database for automatic tuning.

Configure the automatic tuning options ⓘ

Option	Desired state
 FORCE PLAN	☐ ON ☐ OFF INHERIT
 CREATE INDEX	☐ ON ☐ OFF INHERIT
 DROP INDEX	☐ ON ☐ OFF INHERIT

Microsoft

Your company has a SQL Server 2019 instance hosted on an Azure virtual machine. The machine runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

You now have to scale up the virtual machine to 16 vCPUs and 64 GB of memory.

You have to provide the lowest latency for the tempdb database.

You decide to create 2 data files for tempdb

Would this fulfill the requirement?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

As per the Microsoft documentation, you should have 1 file per vCore for TempDB with a limit of 8 files.

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices>

NEW QUESTION: 115

You have the following resources:

- 15 SQL Server on Azure Virtual Machines instances
- 20 Azure SQL databases

You need to recommend a solution to centrally monitor the resources for security vulnerabilities.

What should you include in the recommendation?

A. database audits

B. Microsoft Defender

C. SQL insights

D. Azure SQL Auditing

Answer: B (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/defender-for-cloud/defender-for-sql-introduction>


NEW QUESTION: 116

Hotspot Question

You have an Azure SQL database named DB1 that contains a table named Table1.

You run a query to load data into Table1.

The performance metrics of Table1 during the load operation are shown in the following exhibit.



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

Answer Area

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

- scale the resource
- use an elastic resource
- 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:

Answer Area

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

▼

scale the resource

use an elastic resource

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

Explanation:
You are loading data into Table1, which means there is some inserts going on. Note that insert statements is also considered a query. It should be scale resource in order to reduce the time insert data into Table1.

NEW QUESTION: 117

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.
What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?
A. Configure SalesSQLDb1 as writable.
B. Update the connection strings of SalesSQLDb1App1.
C. Update the firewall rules of SalesSQLDb1.
D. Update the users in SalesSQLDb1.

Answer: (SHOW ANSWER)
Scenario: SalesSQLDb1 uses database firewall rules and contained database users.
Incorrect:
Not C: When using public network access for connecting to the database, we recommend using database-level IP firewall rules for geo-replicated databases.
These rules are replicated with the database, which ensures that all geo-secondaries have the same IP firewall rules as the primary. This approach eliminates the need for customers to manually configure and maintain firewall rules on servers hosting the primary and secondary databases.
Reference:
<https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview#keeping-credentials-and-firewall-rules-in-sync>

NEW QUESTION: 118

You are training a new administrator for your company's Azure data services, which includes database deployed on Azure SQL Database, Azure SQL Managed Instance, and SQL Server on Azure virtual machines (VMs).
You need to identify the fixed roles that are supported by Azure SQL Database only.
Which two fixed roles are available with Azure SQL Database only?
Each correct answer presents a complete solution.
A. db_securityadmin
B. dbcreator
C. loginmanager
D. dbmanager
E. sysadmin

Answer: (SHOW ANSWER)

NEW QUESTION: 119

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 to use SQLDB11 as an elastic job database to run jobs on SQLDB11 and SQLDB22.
What is the minimum number of database scoped credentials required for the elastic jobs?
A. 4

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

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

NEW QUESTION: 120

You have an Azure SQL managed instance named MI1.

You need to implement automatic tuning for the databases of MI1.

What should you do?

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

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

For Azure SQL Managed Instance, the supported option FORCE_LAST_GOOD_PLAN can only be configured through T-SQL. The Azure portal based configuration and automatic index tuning options described in this article do not apply to Azure SQL Managed Instance.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable?view=azuresql>

NEW QUESTION: 121

Hotspot Question

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.

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.	☐	☐

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:

Box 1: Yes

Note: sys.dm_hadr_availability_group_states (Transact-SQL)

Returns a row for each Always On availability group that possesses an availability replica on the local instance of SQL Server. Each row displays the states that define the health of a given availability group.

Box 2: Yes

We see: WHERE primary_replica != @@Servername

The @@Servername variable would be Server1 as the T-SQL statements are run on Server1.

Note: AlwaysOn Availability Groups provide high availability for groups of database by adding secondary replicas. These replicas allow failing over databases in case of a failure. In addition they can be used to offload read workloads or backup tasks.

Box 3: No

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 122

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

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.

Which type of compression provides the greatest space reduction for the database?

- A. page compression
- B. row compression
- C. columnstore compression
- D. columnstore archival compression

Answer: (SHOW ANSWER)

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression. For columnstore tables and indexes, all columnstore tables and indexes always use columnstore compression and this is not user configurable. Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data. You can configure columnstore archival compression.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression>

<https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression?view=sql-server-ver15>

NEW QUESTION: 123

You have an Azure subscription.

You plan to create an instance of SQL Server on Azure Virtual Machines.

You need to ensure that the instance is resilient against datacenter failures.

What should you use?

- A. an availability set
- B. a proximity placement group
- C. an availability zone
- D. Azure Site Recovery

Answer: D (LEAVE A REPLY)

NEW QUESTION: 124

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a Log Analytics workspace named Workspace1.

You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace.

What should you do first?

- A. Create a private endpoint connection on SQLMI1.
- B. Configure Azure SQL Analytics to use Workspace1.
- C. Modify the Compute + storage settings for SQLMI1.
- D. Modify the diagnostic settings for SQLMI1.

Answer: D (LEAVE A REPLY)

Diagnostic settings available for databases in your Azure SQL Managed Instance include:

log: SQL Insights, Query Store Runtime Statistics, Query Store Wait Statistics, and Errors destination details: Send to Log Analytics workspace, Archive to a storage account, Stream to an event hub, Send to partner solution

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/monitoring-sql-managed-instance-azure-monitor?view=azuresql#collection-and-routing>

NEW QUESTION: 125

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements

Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

Hotspot Question

You need to recommend a service tier and a method to offload analytical workloads for the databases migrated from SVR1. The solution must meet the availability and business requirements.

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

NOTE: Each correct selection is worth one point.

Answer Area

Service tier: 

Business Critical

General Purpose

Premium

Method: 

A failover group read-only listener

Geo-replicated secondary replicas

Read scale-out

Answer:
Answer Area



Service tier: 

Business Critical

General Purpose

Premium

Method: 

A failover group read-only listener

Geo-replicated secondary replicas

Read scale-out

Explanation:

The read scale-out feature allows you to offload read-only workloads using the compute capacity of one of the read-only replicas, instead of running them on the read-write replica. Read scale-out is always enabled in the Business Critical service tier of SQL Managed Instance, and for Hyperscale databases with at least one secondary replica.
<https://learn.microsoft.com/en-us/azure/azure-sql/database/read-scale-out?view=azuresql>

NEW QUESTION: 126

Drag and Drop Question

You have an on-premises Microsoft SQL Server failover cluster that contains three nodes.
You need to rename a SQL Server instance in the cluster. The solution must minimize downtime for the cluster.
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.

Action	Answer Area
Stop the SQL Server instance.	1
Rename the SQL Server instance.	2
Pause the SQL Server instance.	3
Start the SQL Server instance.	
Rename the host server.	

Answer:

Action	Answer Area
	1 Rename the SQL Server instance.
	2 Stop the SQL Server instance.
Pause the SQL Server instance.	3 Start the SQL Server instance.
Rename the host server.	

NEW QUESTION: 127

SIMULATION

You need to rebuild the index for PK_Address_AddressID in the SalesLT.Address table in db1.

The solution must meet the following requirements:

- Ensure that 30 percent of the index page is reserved for future updates and inserts.
- Minimize page splits.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Note: The easiest way to fix page splits is to set a fill factor value (for instance, 90) when the offending index is rebuilt. If a fill factor is not set, SQL Server uses the server-level default, which is 100 (this default can be changed, however).

Reorganize or rebuild an index in SQL Server Management Studio

Step 1: In Object Explorer, Expand the database [Here db1] that contains the table on which you want to reorganize an index.

Step 2: Expand the Tables folder.

Step 3: Expand the table [Here: SalesLT.Address] on which you want to reorganize an index.

Step 4: Expand the Indexes folder.

Step 5: Right-click the index [Here for PK_Address_AddressID] with the fill factor that you want to specify and select Properties.

5a) Under Select a page, select Options.

5b) In the Fill factor row, enter the fill factor that you want. [Here: 70%]

Question: Ensure that 30 percent of the index page is reserved for future updates and inserts.

5c) Click OK.

Step 6: Right-click the index you want to reorganize [Here for PK_Address_AddressID] and select Reorganize.

Step 7: In the Reorganize Indexes dialog box, verify that the correct index is in the Indexes to be reorganized grid and click OK.

Step 8: Select the Compact large object column data check box to specify that all pages that contain large object (LOB) data are also compacted.

Step 9: Click OK.

Reference:

<https://www.beyondtrust.com/docs/privileged-identity/faqs/reorganize-and-rebuild-indexes-in-database.htm>

<https://learn.microsoft.com/en-us/sql/relational-databases/indexes/specify-fill-factor-for-an-index>

NEW QUESTION: 128

Hotspot Question

You have the projects shown in the following table.

Name	Description
Project1	A rapidly iterating cloud-based solution that integrates multiple software as a service (SaaS) solutions
Project2	The migration of an existing on-premises Microsoft SQL Server 2019 database to Azure
Project3	An app that supports multiple 20-TB databases

You need to recommend an Azure database solution for each project. The solution must meet the following requirements:

- Project2 must minimize malware threats to the underlying operating system that hosts the migrated database.
- Administrative effort must be minimized where possible.

What should you recommend using 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 SQL Managed Instance
- Azure SQL Database
- Azure Cosmos DB
- SQL Server on Azure VM

Project2: ▼

- Azure SQL Managed Instance
- Azure SQL Database
- Azure Cosmos DB
- SQL Server on Azure VM

Project3: ▼

- Azure SQL Managed Instance
- Azure SQL Database
- Azure Cosmos DB
- SQL Server on Azure VM

Answer:

Answer Area

Project1:

Azure SQL Managed Instance

Azure SQL Database

Azure Cosmos DB

SQL Server on Azure VM

Project2:

Azure SQL Managed Instance

Azure SQL Database

Azure Cosmos DB

SQL Server on Azure VM

Project3:

Azure SQL Managed Instance

Azure SQL Database

Azure Cosmos DB

SQL Server on Azure VM

Explanation:

Box 1: Azure Cosmos DB

A rapidly iterating cloud-based solution that integrates multiple software as a service (SaaS) solutions.

For a rapidly iterating, cloud-based SaaS integration solution, Azure Cosmos DB is a strong recommendation. It offers a schemaless, multi-model database that supports flexible data modeling and querying, which is crucial for agile development and integrating diverse SaaS solutions. Its global distribution and automatic scaling capabilities also ensure high availability and performance as your application grows.

Box 2: Azure SQL Managed Instance

The migration of an existing on-premises Microsoft SQL Server 2019 database to Azure.

Project2 must minimize malware threats to the underlying operating system that hosts the migrated database.

Azure SQL Managed Instance minimizes malware threats to the underlying operating system by leveraging a multi-layered security approach, including threat detection, encryption, and network isolation. It achieves this by handling the operating system and underlying infrastructure at the PaaS level, reducing the attack surface accessible to malicious actors.

Box 3: SQL Server on Azure on VM

An app that supports multiple 20-TB databases.

If you need to store more than 16 TB of data, you should consider moving to Azure SQL Database Hyperscale or SQL Server on Azure Virtual Machines.

Reference:

<https://learn.microsoft.com/en-us/azure/architecture/guide/multitenant/service/cosmos-db>

<https://learn.microsoft.com/en-us/azure/azure-sql/database/security-overview>

NEW QUESTION: 129

Drag and Drop Question

You have an Azure key vault named Vault1 and a SQL server on Azure Virtual Machines instance named SQL1. SQL1 hosts a database named DB1.

You need to configure Transparent Data Encryption (TDE) on DB1 to use a key in Vault1.

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

On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.

On SQL1, create a symmetric key.

On SQL1, create a login from the asymmetric key.

On SQL1, create a login from the symmetric key.

Create an Azure AD service principal and grant the service principal permissions for Vault1.

On SQL1, create an asymmetric key.

Answer:



Answer Area

Microsoft

freepdfdumps.com



Actions



Answer Area

On SQL1, create a symmetric key.

On SQL1, create a login from the symmetric key.

Create an Azure AD service principal and grant the service principal permissions for Vault1.

On SQL1, create an asymmetric key.

On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.

On SQL1, create a login from the asymmetric key.

NEW QUESTION: 130

You have an Azure subscription that contains an Azure SQL database named db1.

You need to implement SQL insights for db1.

Which two resources should you create first? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a storage account
- B. a virtual machine
- C. an Azure logic app
- D. an Azure function
- E. a Log Analytics workspace

Answer: B,E (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-insights-enable?view=azuresql>

NEW QUESTION: 131

Hotspot Question

You have an Azure subscription that is linked to a hybrid Azure Active Directory (Azure AD) tenant. The subscription contains an Azure Synapse Analytics SQL pool named Pool1.

You need to recommend an authentication solution for Pool1. The solution must support multi- factor authentication (MFA) and database-level authentication.

Which authentication solution or solutions should you include in the recommendation? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

MFA:

Azure AD authentication

Microsoft SQL Server authentication

Passwordless authentication

Windows authentication

Database-level authentication:

Application roles

Contained database users

Database roles

Microsoft SQL Server logins

Answer:
Answer Area

MFA:

Azure AD authentication

Microsoft SQL Server authentication

Passwordless authentication

Windows authentication

Database-level authentication:

Application roles

Contained database users

Database roles

Microsoft SQL Server logins

Explanation:

Box 1: Azure AD authentication

Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.

Box 2: Contained database users

Azure Active Directory Uses contained database users to authenticate identities at the database level.

Incorrect:

SQL authentication: To connect to dedicated SQL pool (formerly SQL DW), you must provide the following information:

- Fully qualified servername
- Specify SQL authentication
- Username
- Password
- Default database (optional)

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-authentication>

NEW QUESTION: 132

You have the following Azure Data Factory pipelines:

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

Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2.

Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours.

What should you do to schedule the pipelines for execution?

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

Answer: D (LEAVE A REPLY)

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

NEW QUESTION: 133

You have an Azure subscription.

You plan to deploy four SQL Server on Azure Virtual Machines instances to two Azure regions.

Two instances will be deployed to each region. Each region will host a domain-independent availability group.

You need to recommend a high availability, disaster recovery (HA/DR) solution for the planned deployment. The solution must minimize administrative effort.

What should you include in the recommendation?

- A.** a Hyper-V cluster
- B.** a workgroup cluster
- C.** an Active Directory-detached cluster
- D.** a Network Load Balancing (NLB) cluster

Answer: B (LEAVE A REPLY)

SQL Server on Azure Virtual Machines, Configure a workgroup availability group You can create an Active Directory domain-independent cluster with an Always On availability group; this is also known as a workgroup cluster.

Windows Server 2016 introduced a new type of Windows Server failover cluster based on the foundation of the Active Directory-detached cluster: a workgroup cluster. A workgroup cluster allows SQL Server to deploy an availability group on top of a WSFC that doesn't require AD DS.

SQL Server requires the use of certificates for endpoint security, just as the database mirroring scenario requires certificates. This type of an availability group is called a domain independent availability group.

Deploying an availability group with an underlying workgroup cluster supports the following combinations for the nodes that will make up the WSFC:

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-clusterless-workgroup-configure>

<https://learn.microsoft.com/en-us/sql/database-engine/availability-groups/windows/domain-independent-availability-groups>

NEW QUESTION: 134

You have five instances of SQL Server on Azure Virtual Machines.

What should you use?

You need to monitor Microsoft SQL Server performance for all the instances by consolidating metrics into a single graphic display. The solution must minimize administrative effort.

- A.** Azure Monitor

- B. Log Analytics
- C. SQL Insights
- D. Azure SQL Analytics

Answer: A (LEAVE A REPLY)

The Log Analytics agent is on a deprecation path and won't be supported after August 31, 2024.

Any new data centers brought online after January 1 2024 will not support the Log Analytics agent. If you use the Log Analytics agent to ingest data to Azure Monitor, migrate to the new Azure Monitor agent prior to that date.

<https://learn.microsoft.com/en-us/azure/azure-monitor/agents/log-analytics-agent>

NEW QUESTION: 135

You have an Azure subscription. The subscription contains three virtual machines that run Red Hat Enterprise Linux (RHEL).

You plan to implement a highly available deployment of Microsoft SQL Server 2022 on the virtual machines by using failover cluster instances (FCIs).

You need to prepare for the implementation.

What should you do first on each server?

- A. Create the FCI disk and network resources.
- B. Install a named SQL Server instance.
- C. Install the cluster resource manager and the FCI resource agent.
- D. Install the default SQL Server instance.

Answer: C (LEAVE A REPLY)

To prepare for implementing a highly available deployment of Microsoft SQL Server 2022 on the virtual machines using failover cluster instances (FCIs), the first step you should take on each server is install the cluster resource manager and the FCI resource agent. Before setting up a failover cluster instance, you need to ensure that the cluster infrastructure is in place. This involves installing the cluster resource manager (such as Pacemaker on Linux) and the FCI resource agent, which are essential for managing the failover clustering and ensuring high availability.

NEW QUESTION: 136

You have the following Transact-SQL query.

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

Which column returned by the query represents the free space in each file?

- A. ColumnA
- B. ColumnB

C. ColumnC

D. ColumnD

Answer: ([SHOW ANSWER](#))

Example:

Free space for the file in the below query result set will be returned by the FreeSpaceMB column.

```
SELECT DB_NAME() AS DbName,
```

```
name AS FileName,
```

```
type_desc,
```

```
size/128.0 AS CurrentSizeMB,
```

```
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB FROM sys.database_files WHERE type IN (0,1); Reference:
```

<https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 137

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

This is not about an external table.

Instead, in an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files.

Note: You can use the Get Metadata activity to retrieve the metadata of any data in Azure Data Factory or a Synapse pipeline. You can use the output from the Get Metadata activity in conditional expressions to perform validation, or consume the metadata in subsequent activities.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/control-flow-get-metadata-activity>

NEW QUESTION: 138

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements

Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.
- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.

- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

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

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

<https://azure.microsoft.com/en-us/blog/azure-sql-database-point-in-time-restore>

NEW QUESTION: 139

SIMULATION

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:

To connect to an Azure SQL Database using a private IP address on a virtual network via the Azure portal, you'll need to configure a Private Endpoint.

Connect to an Azure SQL server using an Azure Private Endpoint using the Azure portal Create a Private Endpoint:

Step 1: In the Azure portal, search for and select "Private endpoints".

Step 2: Select "+ Create" to create a new Private Endpoint.

Step 3: On the "Basics" tab, provide a name, resource group, and region.

Step 4: On the "Resource" tab, select "Microsoft.Sql/servers" as the resource type and choose your SQL Server. [Select db1] Step 5: On the "Virtual Network" tab, select your virtual. [Select VNET1] Step 6: On the "DNS" tab, review the DNS settings and make any necessary changes. [Skip] Reference:

<https://learn.microsoft.com/en-us/azure/private-link/tutorial-private-endpoint-sql-portal>

NEW QUESTION: 140

Drag and Drop Question

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

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

Actions

Change the Active Directory Admin on TestServer1

Change the server name and related variables in the templates

From the database project, deploy the database schema and permissions

Add IP addresses to the firewall

From the Azure portal, export the Azure Resource Manager templates

From the Azure portal, deploy the templates.

Answer Area

>

<

↑

↓

Answer:

Actions

Change the Active Directory Admin on TestServer1

Add IP addresses to the firewall

Answer Area

From the Azure portal, export the Azure Resource Manager templates

Change the server name and related variables in the templates

From the Azure portal, deploy the templates.

From the database project, deploy the database schema and permissions

NEW QUESTION: 141

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 MI1. The solution must meet the following requirements:

- The connection must be encrypted.
- Network latency must be minimized.

What should you implement?

- A.** virtual network peering
- B.** private endpoints
- C.** service endpoints
- D.** a site-to-site VPN

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 142

You have an on-premises Microsoft SQL Server 2022 instance named SQL1.

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1.

You need to configure a bidirectional disaster recovery solution between SQL1 and SQLMI1.

What should you include in the solution?

- A.** Azure Database Migration Service
- B.** Log Replay Service (LRS)
- C.** a failover group
- D.** the Managed Instance link

Answer: D (LEAVE A REPLY)

Disaster recovery with Managed Instance link - Azure SQL Managed Instance You can configure a hybrid disaster recovery solution between SQL Server hosted anywhere and Azure SQL Managed Instance by using the Managed Instance link.

With SQL Server 2022, either SQL Server or Azure SQL Managed Instance can be the primary and you can establish the link initially from either SQL Server or SQL Managed Instance. You can fail over between SQL Server and Azure SQL Managed Instance in either direction, as needed.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-disaster-recovery>

NEW QUESTION: 143

You have an Azure Synapse Analytics dedicated SQL pool.

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

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

Which statement accurately describes the dbo.FactInternetSales table?

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

Answer: D (LEAVE A REPLY)

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

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

This is a very quick and simple way to see the number of table rows that are stored in each of the

60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.

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

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute>

<https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distributed-data-skew.md>

NEW QUESTION: 144

You have an Azure SQL database named DB1. You run a query while connected to DB1.

You review the actual execution plan for the query, and you add an index to a table referenced by the query.

You need to compare the previous actual execution plan for the query to the Live Query Statistics.

What should you do first in Microsoft SQL Server Management Studio (SSMS)?

A. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.

B. Run the SET SHOWPLAN_ALL Transact-SQL statement.

C. Save the actual execution plan.

D. Enable Query Store for DB1.

Answer: C (LEAVE A REPLY)

The Plan Comparison menu option allows side-by-side comparison of two different execution plans, for easier identification of similarities and changes that explain the different behaviors for all the reasons stated above. This option can compare between:

Two previously saved execution plan files (.sqlplan extension).

One active execution plan and one previously saved query execution plan.

Two selected query plans in Query Store.

NEW QUESTION: 145

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: B,C (LEAVE A REPLY)

Auditing an instance of SQL Server or a SQL Server database involves tracking and logging events that occur on the system. The SQL Server Audit object collects a single instance of server-level or database-level actions and groups of actions to monitor. The audit is at the SQL Server instance level. You can have multiple audits per SQL Server instance. The Database- Level Audit Specification object belongs to an audit. You can create one database audit specification per SQL Server database per audit.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/create-a-server-audit-and-database-audit-specification>

NEW QUESTION: 146

Hotspot Question

You have an on-premises Microsoft SQL Server instance named SQLSVR1 that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL database named SQLDB1.

You need to migrate the data stored in DB1 to SQLDB1 by using SQL Server replication. The solution must minimize the performance impact on DB1.

How should you configure the replication? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



Configure SQLDB 1:

As a push subscriber only

As a push subscriber and a distributor

To use a pull subscription only

Type of replication:

Snapshot

Transactional with peer-to-peer replication

Transactional with updatable subscriptions

Answer:

Answer Area

Configure SQLDB 1:

As a push subscriber only

As a push subscriber and a distributor

To use a pull subscription only

Type of replication:

Snapshot

Transactional with peer-to-peer replication

Transactional with updatable subscriptions

Explanation:

Box 1: As a push subscriber and a distributor

To migrate a SQL Server database to Azure SQL Database using replication, you'll need to set up transactional replication with the on-premises SQL Server as the publisher and distributor and the Azure SQL Database as the subscriber. This process involves configuring the replication topology, including the publication, subscription, and the distributor. You'll also need to ensure proper authentication and firewall rules are set up.

Box 2: Transactional with updatable subscriptions

Snapshot and one-way transactional replication are supported. Peer-to-peer transactional replication and merge replication aren't supported.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/replication-to-sql-database>

NEW QUESTION: 147

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

- A. Input Deserialization Errors
- B. Late Input Events
- C. Early Input Events
- D. Watermark delay

Answer: D (LEAVE A REPLY)

The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.

The watermark delay metric can rise due to:

1. Not enough processing resources in Stream Analytics to handle the volume of input events.
2. Not enough throughput within the input event brokers, so they are throttled.
3. Output sinks are not provisioned with enough capacity, so they are throttled.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling>

NEW QUESTION: 148

Hotspot Question

You have an Azure SQL database named db1.

You need to retrieve the resource usage of db1 from the last week.

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

NOTE: Each correct selection is worth one point.

Answer Area

SELECT *

FROM

▼

sys.dm_db_resource_stats
sys.dm_exec_requests
sys.dm_user_db_resource_governance
sys.resource_stats

WHERE database_name = 'db1' AND

start_time >

▼

DATEADD
DATEDIFF
DATEPART
TOTIMEOFFSET

(day, -7, GETDATE())

Microsoft

ORDER BY start_time DESC;

Answer:

Answer Area

SELECT *

FROM

▼

sys.dm_db_resource_stats

sys.dm_exec_requests

sys.dm_user_db_resource_governance

sys.resource_stats

WHERE database_name = 'db1' AND

start_time >

▼

DATEADD

DATEDIFF

DATEPART

TODATETIMEOFFSET

(day, -7, GETDATE())

ORDER BY start_time DESC

Explanation:

Box 1: sys.resource_stats

sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.

Box 2: DateAdd

The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

DECLARE @s datetime;

DECLARE @e datetime;

SET @s= DateAdd(d,-7,GetUTCDate());

SET @e= GETUTCDATE();

SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization FROM sys.resource_stats WHERE start_time BETWEEN @s AND @e GROUP BY database_name HAVING

AVG(avg_cpu_percent) >= 80 Incorrect Answers:

sys.dm_exec_requests:

sys.dm_exec_requests returns information about each request that is executing in SQL Server. It does not have a column named database_name.

sys.dm_db_resource_stats:

sys.dm_db_resource_stats does not have any start_time column.

Note: sys.dm_db_resource_stats returns CPU, I/O, and memory consumption for an Azure SQL Database database. One row exists for every 15 seconds, even if there is no activity in the database. Historical data is maintained for approximately one hour.

Sys.dm_user_db_resource_governance returns actual configuration and capacity settings used by resource governance mechanisms in the current database or elastic pool. It does not have any start_time column.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-database>

NEW QUESTION: 149

Hotspot Question

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



CREATE USER "AutomationAccount"

FROM

ASYMETRIC KEY
EXTERNAL PROVIDER
FOR CERTIFICATE

WITH

KEY
LOGIN
OBJECT_ID

 = '1133a5a2-4c19-47d7-c2d5-42b289997918f'

EXEC sp_addrolemember 'db_owner', "AutomationAccount"

Answer:

Answer Area



CREATE USER "AutomationAccount"

FROM

ASYMETRIC KEY
EXTERNAL PROVIDER
FOR CERTIFICATE

WITH

KEY
LOGIN
OBJECT_ID

 = '1133a5a2-4c19-47d7-c2d5-42b289997918f'

EXEC sp_addrolemember 'db_owner', "AutomationAccount"

Explanation:

Box 1: FROM EXTERNAL PROVIDER

Create a user-assigned managed identity

Box 2: OBJECT_ID

The above syntax allows creation of Microsoft Entra users without validation. For this to work, the Object Id of the Microsoft Entra principal would have to be supplied and used as an SID in the T- SQL statement.

The validity check of the Object Id is the responsibility of the user running the T-SQL statement.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-azure-ad-user-assigned-managed-identity>

NEW QUESTION: 150

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

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

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

- A. SQLAgentOperatorRole
- B. SQLAgentReaderRole
- C. SQLAgentUserRole
- D. db_owner

Answer: A (LEAVE A REPLY)

NEW QUESTION: 151

Hotspot Question

You have an Azure SQL database named db1 on a server named server1.

You use Query Performance Insight to monitor db1.

You need to modify the Query Store configuration to ensure that performance monitoring data is available as soon as possible.

Which configuration setting should you modify and which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS

INTERVAL_LENGTH_MINUTES

MAX_PLANS_PER_QUERY

QUERY_CAPTURE_MODE

Value:

1

60

CUSTOM

ON



Answer:

Answer Area



Configuration setting:

DATA_FLUSH_INTERVAL_SECONDS

INTERVAL_LENGTH_MINUTES

MAX_PLANS_PER_QUERY

QUERY_CAPTURE_MODE

Value:

1

60

CUSTOM

ON

Explanation:

Box 1: INTERVAL_LENGTH_MINUTES

INTERVAL_LENGTH_MINUTES defines size of time window during which collected runtime statistics for query plans are aggregated and persisted. Every active query plan has at most one row for a period of time defined with this configuration.

Default: 60

Box 2: 1

Statistics Collection Interval (INTERVAL_LENGTH_MINUTES): Defines the level of granularity for the collected runtime statistic, expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues.

Use SQL Server Management Studio or Transact-SQL to set a different value for Statistics Collection Interval:

```
ALTER DATABASE [QueryStoreDB]
SET QUERY_STORE (INTERVAL_LENGTH_MINUTES = 60);
```

Reference:

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

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 152

Drag and Drop Question

You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named App1. App1 uses db1.

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

- App1 must be able to consume all available CPU resources.
- App1 must have at least half of the available CPU resources always

available.

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

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

ACTIONS

Create a plan.

Create a classifier function in db1.

Create a workload group.

Create a classifier function in the master database.

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

Answer Area

Answer:

ACTIONS

Create a plan.

Create a classifier function in db1.

Answer Area

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

Create a workload group.

Create a classifier function in the master database.

NEW QUESTION: 153

Drag and Drop Question

You have a Microsoft SQL Server instance that hosts a database named DB1. DB1 contains a table named Sales.Customers.

You need to implement row-level security (RLS) to meet the following requirements:

- Customers must only see data for the customer category to which they are assigned.
- Employees must be able to see all the data in the Sales.Customers

table.

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
To the Sales.Customers table, add a check constraint that restricts data visibility by category.	1
Create a database role. Grant the SELECT permission for the Sales.Customers table the role.	2
Create an inline table-valued function that filters rows based on the customer category.	3
Create a security policy that adds a filter predicate to the Sales.Customers table.	
To the security policy, add a security predicate that invokes the inline table-valued function.	
Grant the SELECT permission for the inline table-valued function to all users.	



Answer:

Actions	Answer Area
To the Sales.Customers table, add a check constraint that restricts data visibility by category.	1 Create a database role. Grant the SELECT permission for the Sales.Customers table the role.
Create a database role. Grant the SELECT permission for the Sales.Customers table the role.	2 Create an inline table-valued function that filters rows based on the customer category.
Create an inline table-valued function that filters rows based on the customer category.	3 Create a security policy that adds a filter predicate to the Sales.Customers table.
Create a security policy that adds a filter predicate to the Sales.Customers table.	
To the security policy, add a security predicate that invokes the inline table-valued function.	
Grant the SELECT permission for the inline table-valued function to all users.	

NEW QUESTION: 154

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.
- * 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.
- * 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.

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.



Role Assignments

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



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: D (LEAVE A REPLY)

Scenario: Exhibit show idle periods with intermittent spikes.

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns.

Incorrect Answers:

B: Hyperscale is for large databases

Reference:

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

NEW QUESTION: 155

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

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

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

NOTE: Each correct selection is worth one point.


- A. Create a date dimension table that has a DateTime key.
- B. Create a date dimension table that has an integer key in the format of YYYYMMDD.

- C. Use built-in SQL functions to extract date attributes.
- D. Use integer columns for the date fields.
- E. Use DateTime columns for the date fields.

Answer: B,D (LEAVE A REPLY)

Why use a Date Dimension Table in a Data Warehouse.

The Date dimension is one of these dimension tables related to the Fact. Here is a simple Data Diagram for a Data Mart of Internet Sales information for the Adventure Works DW database which can be obtained for free from CodePlex or other online sources.



The relationship is created by the surrogate keys columns (integer data type) rather than the date data type.

The query users have to write against a Data Mart are much simpler than against a transaction database. There are less joins because of the one to many relationships between the fact dimension table(s). The dimension tables are confusing to someone who has been normalizing databases as a career. The dimension is a flattened or de-normalized table. This creates cases of duplicate data, but the simplistic query overrides the duplicate data in a dimensional model.

Reference:

<https://www.mssqltips.com/sqlservertip/3117/defining-role-playing-dimensions-for-sql-server-analysis-services/>

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

NEW QUESTION: 156

You have an Azure SQL database named sqlldb1.

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

What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

Answer: C (LEAVE A REPLY)

To avoid Query Store transitioning to read-only state, the runtime statistics for the Last Stale Query Threshold days must fit in the defined Max Size (MB). If the Statistics Collection Interval is increased then less space will be used for the persisted runtime statistics and there will be more free space for new data. The Data Flush Interval can delay the transitioning to read-only state, but not minimize the possibility it occurs.

Reference:

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

NEW QUESTION: 157

Drag and Drop Question

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 SOL 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

Assign the proxy to the job step.

Create a login.

Create a database user.

Create a credential.

Create a proxy.

Answer Area

➤

➤

⬅

⬆

⬆

⬆

Answer:

Actions

Create a login.

Create a database user.

Answer Area

Create a credential.

Create a proxy.

Assign the proxy to the job step.

➤

➤

⬅

⬆

⬆

⬆

NEW QUESTION: 158

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 another Azure SQL database in the Hyperscale service tier.
- B.** Export SQLDB1.
- C.** Modify the Backup storage redundancy setting for SQLDB1.
- D.** Create a logical server for Azure SQL Database.

Answer: D (LEAVE A REPLY)

Configure active geo-replication and failover (Azure SQL Database)

Active geo-replication is configured per database.

Step 1: Add a secondary database

The secondary database has the same name as the primary database and has, by default, the same service tier and compute size. The secondary database can be a single database or a pooled database.

After the secondary is created and seeded, data begins replicating from the primary database to the new secondary database.

If your secondary replica is used only for disaster recovery (DR), and doesn't have any read or write workloads, you can save on licensing costs by designating the database for standby when you configure a new active geo-replication relationship.

(Step 2: Initiate a failover)

The secondary database can be switched to become the primary.

The command immediately switches the secondary database into the primary role. This process normally should complete within 30 seconds or less.

Both databases are unavailable, for up to 25 seconds, while the roles are switched. If the primary database has multiple secondary databases, the command automatically reconfigures the other secondaries to connect to the new primary. The entire operation should take less than a minute to complete under normal circumstances.

Reference:

<https://docs.azure.cn/en-us/azure-sql/database/active-geo-replication-configure-portal>

NEW QUESTION: 159

Hotspot Question

You have an Azure subscription.

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

You need to configure the resource section of the template to ensure that the deployment always maximizes performance for network connections to the server, regardless of whether the connection originates from within or outside of Azure.

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

NOTE: Each correct selection is worth one point.

```



resource server 'Microsoft.Sql/servers@2022-05-01-preview' = {
  name: serverName
  location: location
  properties: {
    administratorLogin: administratorLogin
    administratorLoginPassword: administratorLoginPassword
    version: '12.0'
  }
}

resource connectionPolicy 'Microsoft.Sql/servers/connectionPolicies@2022-05-01-preview' =
  parent:
    Microsoft.Sql/servers
    server
    serverName
  name: 'connectionPolicyName'
  properties: {
    connectionType:
      'Default'
      'Proxy'
      'Redirect'
  }
}

```

Answer:

Answer Area

```
resource server 'Microsoft.Sql/servers@2022-05-01-preview' = {  
  name: serverName  
  location: location  
  properties: {  
    administratorLogin: administratorLogin  
    administratorLoginPassword: administratorLoginPassword  
    version: '12.0'  
  }  
}  
  
resource connectionPolicy 'Microsoft.Sql/servers/connectionPolicies@2022-05-01-preview' = {  
  parent:   
  name: 'connectionPolicyName'  
  properties: {  
    connectionType:   
  }  
}
```

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.** Deny public access.
- B.** Create a private endpoint.
- C.** Create a network security group (NSG).
- D.** Create a server-level firewall rule.
- E.** Create a database-level firewall rule.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 161

You have an Azure SQL Database elastic pool that contains 10 databases.

You receive the following alert.

Msg 1132, Level 16, state 1, Line 1

The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.

You need to resolve the alert. The solution must minimize administrative effort.

Which three actions can you perform? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: **A,C,E** ([LEAVE A REPLY](#))

<https://learn.microsoft.com/en-us/azure/azure-sql/database/file-space-manage?view=azuresql-db>

NEW QUESTION: 162

Hotspot Question

You have an Azure SQL database named DB1.

You need to create a second instance of DB1 in a different Azure region. The solution must meet the following requirements:

- Ensure that the instance is readable and available to handle incoming queries in the event of a local Azure region failure.
- Minimize administrative effort.
- Minimize costs.

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

NOTE: Each correct selection is worth one point.

Answer Area



Microsoft

Create the instance of DB1 by using:

	▼
A database copy	
Geo-replication	
Failover groups	
Zone redundancy	

Type of database to create:

	▼
A non-readable replica	
A readable replica	
A snapshot	
An active secondary replica	

Answer:

Answer Area

 Microsoft

Create the instance of DB1 by using:

- A database copy
- Geo-replication**
- Failover groups
- Zone redundancy

Type of database to create:

- A non-readable replica
- A readable replica
- A snapshot
- An active secondary replica**

Explanation:

Box 1: Geo-replication

To create a readable secondary instance of an Azure SQL Database in a different region while ensuring high availability, low administration, and low costs, you should use Active Geo- Replication.

Active Geo-Replication is the optimal built-in feature for this scenario. It creates a readable, secondary database in a different geographic region. It handles asynchronous replication of data automatically.

Box 2: An active secondary replica

You should create an active secondary replica using Azure SQL Database geo-replication to meet the requirements. This solution directly satisfies your needs for readability, high availability during a regional failure, low administration, and cost-efficiency.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>


NEW QUESTION: 163

Drag and Drop Question


You have an Azure SQL database named DB1.

You need to create a partitioned table in 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.



Answer:



Explanation:

<https://learn.microsoft.com/en-us/sql/relational-databases/partitions/partitioned-tables-and-indexes?view=sql-server-ver16>

NEW QUESTION: 164

You have an Azure subscription that contains three instances of SQL Server on Azure Virtual Machines.

You plan to implement a disaster recovery solution.

You need to be able to perform disaster recovery drills regularly. The solution must meet the following requirements:

- Minimize administrative effort for the recovery drills.
- Isolate the recovery environment from the production environment

What should you use?

- A. Azure Backup
- B. Azure Site Recovery
- C. Recovery Services vaults
- D. native Microsoft SQL Server backup

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

NEW QUESTION: 165

Hotspot Question

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.

Answer Area

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:

Answer Area

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:
Box 1: No
Just one failure within the 5-minute interval.
Box 2: No
Just two failures within the 5-minute interval.
Box 3: No
Just two failures within the 5-minute interval.
Reference:
<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

NEW QUESTION: 166

You have an Azure subscription that contains an Azure SQL database named DB1. DB1 contains a table named Table1.

You need to recommend a security solution that meets the following requirements:

- Ensures that different departments at your company can query Table1
- Ensures that each department can view only the rows in Table1 that are associated to that department
- Ensures that the solution can scale as new departments are added
- Minimizes administrative effort

What should you include in the recommendation?

- A. Encrypt the department columns in Table1 by using Always Encrypted.
- B. Enable row-level security (RLS) and add a security predicate function that filters by department.
- C. Create database roles. Grant the roles the SELECT permission for the filtered views of each department.
- D. Implement dynamic data masking on the columns in Table1 that contain sensitive data.

Answer: B (LEAVE A REPLY)

To implement this efficiently while minimizing administrative overhead as you scale, you should combine Row-Level Security (RLS) with Microsoft Entra groups.

Why this scales:

Zero SQL Changes: When a new department is added, you simply create a new Entra group and add a new row to your data with that group name. You do not need to update the SQL function or policy.

Centralized Management: Adding or removing employees happens in Microsoft Entra ID, automatically updating their access in SQL without DBA intervention.

Performance: The predicate is evaluated at the storage engine level, ensuring users only "see" the rows they have access to, which also prevents accidental data leakage in logs or execution plans.

Note:

Here is the most scalable way to structure this:

1. Group Strategy

Instead of mapping individual users to departments in your SQL code, create a Microsoft Entra ID Security Group for each department. Users should be added to these groups at the identity provider level, not within the database.

2. The Filter Predicate

Create a logic function that checks the user's group membership. The most efficient method is using the IS_MEMBER() function or by mapping the SESSION_USER to a mapping table.

```
CREATE FUNCTION Security.fn_departmentSecurityPredicate(@Department AS sysname) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS fn_security_predicate_result WHERE
-- Allows a user to see rows where their Entra Group matches the department column IS_MEMBER(@Department) = 1 OR IS_MEMBER('DatabaseAdmins') = 1; -- Example exception for IT/Admins
```

3. Apply the Security Policy

Apply the function to your target table. By using STATE = ON, the filter is automatically applied to all incoming queries.

```
CREATE SECURITY POLICY DepartmentFilter
ADD FILTER PREDICATE Security.fn_departmentSecurityPredicate(DepartmentColumn) ON dbo.YourTableName WITH (STATE = ON); Reference:
https://medium.com/@aakanksha.s0905/implementing-row-level-security-rls-with-azure-active-directory-aad-groups-in-sql-views-52479135d91a
```

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 167

Hotspot Question

You plan to deploy an instance of SQL Server on Linux Azure Virtual Machines. The instance will run Microsoft SQL Server 2022 and use the SQL Server IaaS Agent extension for Linux. Which Linux operating system should you deploy, and which benefit will the SQL Server IaaS Agent extension provide? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Operating system:

Red Hat Enterprise Linux (RHEL)

SUSE Linux Enterprise Server (SLES)

Ubuntu Linux

Benefit:

Automatic backup

Automatic patching

Registration of a SQL virtual machine resource in Azure

Answer:

Answer Area

Operating system:

Red Hat Enterprise Linux (RHEL)
SUSE Linux Enterprise Server (SLES)
Ubuntu Linux

Benefit:

Automatic backup
Automatic patching
Registration of a SQL virtual machine resource in Azure

Explanation:

Operating System: Ubuntu Linux

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/linux/sql-server-iaas-agent-extension-linux?view=azuresql&tabs=azure-powershell#limitations> Limitations: The Linux SQL IaaS Agent extension has the following limitations:

Only SQL Server VMs running on the Ubuntu Linux operating system are supported. Other Linux distributions are not currently supported.

Benefit: Registration of a SQL virtual machine resource in Azure

In below link, automated backups and patching are NOT present in the Benefits list for Linux:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/linux/sql-server-iaas-agent-extension-linux?view=azuresql&tabs=azure-powershell>

NEW QUESTION: 168

You provision an Azure SQL Managed Instance database named MyDevData. The database will be used by in-house development for application development projects.

DevSupport custom database role members will use dynamic management views (DMVs) to retrieve performance and health information about MyDevData.

You need to ensure that DevSupport members can view information through DMVs.

Which statement should you use?

Each correct answer presents part of the solution.

- A. GRANT VIEW SERVER STATE TO DevSupport
- B. GRANT VIEW REFERENCES TO DevSupport
- C. GRANT VIEW DATABASE STATE TO DevSupport
- D. GRANT VIEW DEFINITION TO DevSupport

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

NEW QUESTION: 169

You have SQL Server on an Azure virtual machine.

You need to add a 4-TB volume that meets the following requirements:

- Maximizes IOPs
- Uses premium solid state drives (SSDs)

What should you do?

- A. Attach two mirrored 4-TB SSDs.
- B. Attach a stripe set that contains four 1-TB SSDs.
- C. Attach a RAID-5 array that contains five 1-TB SSDs.
- D. Attach a single 4-TB SSD.

Answer: (SHOW ANSWER)

For more throughput, you can add additional data disks and use disk striping.

Reference:

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

NEW QUESTION: 170

Drag and Drop Question

You have an Azure subscription that contains an Azure SQL database named DB1.

You plan to perform a classification scan of DB1 by using Azure Purview.

You need to ensure that you can register DB1.

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 Purview Studio, create a scan trigger.

Modify the Access control (IAM) settings for DB1.

Create a system-assigned managed identity.

Create a Purview account.

From the Purview Studio, create a scan rule set.

Answer Area

Answer:

Actions

From the Purview Studio, create a scan trigger.

From the Purview Studio, create a scan rule set.

Answer Area

Create a Purview account.

Modify the Access control (IAM) settings for DB1.

Create a system-assigned managed identity.

NEW QUESTION: 171

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. Update-SqlVM
- B. Set-AzSqlServerConfigurationOption
- C. Set-AzSqlDatabase
- D. Set-AzSqlInstance

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

To configure an Azure SQL Managed Instance to use the Business Critical service tier with PowerShell, you'll use the Set-AzSqlInstance cmdlet. This cmdlet allows you to modify various properties of an existing managed instance, including its service tier. You'll need to specify the resource group, instance name, and the desired service tier as "BusinessCritical".

Example:

```
$resourceGroupName = "YourResourceGroupName"
```

```
$managedInstanceName = "YourManagedInstanceName"
```

```
$serviceTier = "BusinessCritical"
```

```
Set-AzSqlInstance -ResourceGroupName $resourceGroupName -Name
```

```
$managedInstanceName -ServiceTier $serviceTier
```

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/scripts/create-configure-managed-instance-powershell>

NEW QUESTION: 172

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

Customer contains a column named Forename. Forename contains the first name of each customer. Forename is masked in the database by using dynamic data masking.

From the Azure portal, you review classification recommendations and discover that Forename does NOT have a Confidential sensitivity label.

What is the cause of the incorrect classification?

A. The classification engine did not run yet.

B. The database is using Always Encrypted.

C. The data in Forename is masked already.

D. The classification is based on the column name.

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

The fact that the classification engine has not run yet is the most likely reason why the column lacks a Confidential sensitivity label.

In Azure SQL Database, Dynamic Data Masking (DDM) and Data Discovery and Classification are two entirely independent features. Applying a mask to a column does not automatically assign a sensitivity label to it; the classification engine must scan the database schema to discover and recommend those labels.

Independent Features: Dynamic Data Masking obfuscates real-time data for non-privileged users.

Data Discovery and Classification tags data based on its business sensitivity. They do not automatically trigger each other.

Scan Requirement: The automated classification engine relies on regular schema scans to evaluate column names (like first_name) against built-in or custom information types. If it hasn't run, no recommendations will exist.

Manual Step Needed: Even after the engine runs, recommendations must be explicitly accepted and saved by an administrator before the labels are permanently applied to the columns.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview>

NEW QUESTION: 173

Hotspot Question

You have an Azure subscription that contains an app named App and an Azure SQL database named DB1. App1 queries DB1.

You use Query Performance Insight to get query details for DB1 as shown in the exhibit. (Click the Exhibit tab.)

QUERY ID	CPU[%]	↑↓	DATA IO[%]	↑↓	LOG IO[%]	↑↓	DURATION(h:mm:ss)	↑↓	EXECUTIONS COUNT
35563	0.43		0		0		00:01:20.730		17035
35591	0.4		0		0		00:02:59.340		15086
35541	0.31		0		0		00:11:05.440		7268
35602	0.28		0		0		00:20:50.170		99
35612	0.12		0		0		00:05:24.410		7

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

Query ID [answer choice] is the least performant query overall.

35602

35612

35563

35591

Query ID [answer choice] consumed the most database resources.

35602

35612

35563

35591

Answer:
Answer Area

Query ID [answer choice] is the least performant query overall.

35602

35612

35563

35591

Query ID [answer choice] consumed the most database resources.

35602

35612

35563

35591

Explanation:
Box 1: 35591
Query 35591 has the longest duration.
Review top queries per duration
Two metrics in Query Performance Insight can help you find potential bottlenecks: duration and execution count. Long-running queries have the greatest potential for locking resources longer, blocking other

users, and limiting scalability. They're also the best candidates for optimization.

Box 2: 35563

Query 35563 has the highest execution count.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/query-performance-insight-use>

NEW QUESTION: 174

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: D (LEAVE A REPLY)

Elastic jobs support two authentication methods to connect to target databases: Microsoft Entra ID (Azure AD) authentication using a user-assigned managed identity (UMI) or database-scoped credentials.

Microsoft recommends using user-assigned managed identity (UMI) for authentication because it integrates with Azure AD and supports Conditional Access policies, including location-based policies.

NEW QUESTION: 175

SIMULATION

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:

To create an email alert for high CPU percentage utilization in an Azure database within the Azure portal, you'll first need to navigate to your database resource, then to the "Alerts" section.

From there, you'll create a new alert rule, define the CPU percentage threshold, configure an action group for email notifications, and finally, create the alert rule.

Create alerts for Azure SQL Database and Azure Synapse Analytics using the Azure portal Step 1: In the Azure portal, locate the resource you are interested in monitoring and select it.

[Select db1]

Step 2: In the resource menu under Monitoring, select Alerts. The text and icon might vary slightly for different resources.

Step 3: Select the + Create button, then Alert rule.

Step 4: On the Create an alert rule page, the Scope is automatically configured to the individual resource.

Capture activity peaks or sustained resource stress with alerts on the Metrics signal category.

Step 5: On the Condition tab, select See all signals and Select a signal from the list of Metrics.

Step 6: Select the desired metric, for example CPU percentage. Select Apply. [Select CPU percentage] Step 7: After you select a signal, the Alert logic options appear. A preview of recent activity in this resource for that signal is also displayed.

Step 8: Configure a Threshold type to determine when the alert will take action. Choose Aggregation type, Operator, and Threshold value as desired. A typical threshold is: Static, Maximum, Greater than, 80%.

Threshold: [Select 50%]

Operator: [Select Greater than]

Aggregation granularity: [Select 5 minutes]

Select the interval that's used to group the data points by using the aggregation type function.

Choose an Aggregation granularity (period) that's greater than the Frequency of evaluation to reduce the likelihood of missing the first evaluation period of an added time series.

Frequency of evaluation: [Select one minute]

Select how often the alert rule is to be run. Select a frequency that's smaller than the aggregation granularity to generate a sliding window for the evaluation.

Step 9: Under When to evaluate, determine the desired frequency of evaluation. Use the Check every and Lookback period dropdown lists.

Step 10: Optionally, you can add multiple conditions for this alert, choose the Add condition alert if desired.

Step 11: Select Next: Actions >.

Create the new alert rule

Step 12: Use the Use quick actions feature to create a new action group and provide an email or Azure mobile app notification target.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/alerts-create?>

NEW QUESTION: 176

Drag and Drop Question

You have SQL Server on an Azure virtual machine named SQL1.

SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator.

You need to ensure that dbadmin1 receives an email alert if a job fails.

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

Actions

Create a job alert

Create a job notification

Enable Database Mail

Enable the email settings for the SQL Server Agent

Create a job target

Answer Area


→

←

↑

↓

Answer:



Explanation:

Step 1: Enable Database Mail

Database Mail must be enabled.

Step 2: Enable the email settings for the SQL Server Agent.

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Step 3: Create a job notification

Example:

-- adds an e-mail notification for the specified alert (Test Alert)

-- This example assumes that Test Alert already exists

-- and that François Ajenstat is a valid operator name.

USE msdb ;

GO -

EXEC dbo.sp_add_notification -

@alert_name = N'Test Alert',

@operator_name = N'François Ajenstat',

@notification_method = 1 ;

GO -

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status>

<https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator>

NEW QUESTION: 177

Hotspot Question

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:

Answer Area

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: Azure SQL Managed Instance
Managed Instance will support native BACKUP like SQL Server.
COPY_ONLY Backup
Once you create a credential, you can backup any database using standard BACKUP T-SQL command:
BACKUP DATABASE tpcc2501
TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak' WITH COPY_ONLY The only mandatory property is COPY_ONLY. Azure SQL Managed Instance owns the backup chain. This means that you cannot perform database backup that breaks the backup chain. Azure SQL Managed Instance allows you to perform only COPY_ONLY backups.
Box 2: Auto-failover group

For geographic failover of instances of SQL Managed Instance, use Auto-failover groups.

The auto-failover groups feature allows you to manage the replication and failover of some or all databases on a logical server to another region.

Reference:

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

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

NEW QUESTION: 178

Case Study 5 - ADatum Corporation

Overview

ADatum Corporation is a financial services company that has a main office in New York City.

Existing Environment

Licensing Agreement

ADatum has a Microsoft Volume Licensing agreement that includes Software Assurance.

Network Infrastructure

ADatum has an on-premises datacenter and an Azure subscription named Sub1.

Sub1 contains a virtual network named Network1 in the East US Azure region.

The datacenter is connected to Network1 by using a Site-to-Site (S2S) VPN.

Identity Environment

The on-premises network contains an Active Directory Domain Services (AD DS) forest.

The forest contains a single domain named corp.adatum.com.

The corp.adatum.com domain syncs with a Microsoft Entra tenant named adatum.com.

Database Environment

The datacenter contains the servers shown in the following table.

Name	Operating system	Description
SVR1	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR2	Windows Server 2016	Has Microsoft SQL Server 2016 Enterprise installed. Has an Always On availability group named AG1 that contains two databases named DB1 and DB2.
SVR3	Windows Server 2019	Has Microsoft SQL Server 2019 Enterprise installed that contains a database named DB3.

DB1 and DB2 are used for transactional and analytical workloads by an application named App1.

App1 runs on Microsoft Entra hybrid joined servers that run Windows Server 2022. App1 uses Kerberos authentication.

DB3 stores compliance data used by two applications named App2 and App3.

DB3 performance is monitored by using Extended Events sessions, with the event_file target set to a file share on a local disk of SVR3.

Resource allocation for DB3 is managed by using Resource Governor.

Requirements

Planned Changes

ADatum plans to implement the following changes:

- Deploy an Azure SQL managed instance named Instance1 to Network1.
- Migrate DB1 and DB2 to Instance1.
- Migrate DB3 to Azure SQL Database.

- Following the migration of DB1 and DB2, hand over database development to remote developers who use Microsoft Entra joined Windows 11 devices.
- Following the migration of DB3, configure the database to be part of an auto-failover group.

Availability Requirements

ADatum identifies the following post-migration availability requirements:

- For DB1 and DB2, offload analytical workloads to a read-only database replica in the same Azure region.
- Ensure that if a regional disaster occurs, DB1 and DB2 can be recovered from backups.
- After the migration, App1 must maintain access to DB1 and DB2.
- For DB3, manage potential performance issues caused by resource demand changes by App2 and App3.
- Ensure that DB3 will still be accessible following a planned failover.
- Ensure that DB3 can be restored if the logical server is deleted.
- Minimize downtime during the migration of DB1 and DB2.

Security Requirements

ADatum identifies the following security requirements for after the migration:

- Ensure that only designated developers who use Microsoft Entra joined Windows 11 devices can access DB1 and DB2 remotely.
- Ensure that all changes to DB3, including ones within individual transactions, are audited and recorded.

Management Requirements

ADatum identifies the following post-migration management requirements:

- Continue using Extended Events to monitor DB3.
- In Azure SQL Database, automate the management of DB3 by using elastic jobs that have database-scoped credentials.

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs whenever possible, without affecting other requirements.
- Minimize administrative effort.

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 a private endpoint for connectivity to DB3.
- B.** Configure Microsoft Entra authentication for the logical server that hosts DB3.
- C.** Create a database that has database-scoped credentials.
- D.** Create database-scoped credentials in DB3.

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

NEW QUESTION: 179

You plan to deploy an Azure SQL managed instance.

You need to restore database backups across regions.

Which type of storage account should you use?

- A.** locally-redundant storage (LRS)
- B.** geo-redundant storage (GRS)
- C.** zone-redundant storage (ZRS)
- D.** geo-zone-redundant storage (GZRS)

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

NEW QUESTION: 180

Your company has an Azure SQL Database instance. This database is going to be added to a failover group. It has to be ensured that the database becomes available as quickly as possible if a major outage occurs. The Read/Write failover policy is set to Automatic.

Which of the following can occur as a result of this configuration? (Select TWO)

- A. If a datacenter or Azure regional outage does occur , then the databases can fail over automatically.
- B. If an outage does occur , then the databases in the primary instance can fail over immediately.
- C. If an outage does occur, you can selectively failover databases
- D. If an outage does occur, you can set a different grace period to fail over each database.

Answer: A,D (LEAVE A REPLY)

If you set the Read/Write failover policy as Automatic, in the event of a datacenter or Azure regional outage, the databases will failover automatically.

For Automatic failover, you can also set the Read/Write grace period as shown below.

Failover group

 Create a failover group to automatically failover databases in it.

Failover group name *

Enter failover group name...

.database.windows.net

 Microsoft

*Secondary server

Configure required settings

Read/Write failover policy

Automatic

Read/Write grace period (hours)

1 hours

Database within the group

Select databases to add

Option B is incorrect because this is based on a datacenter or Azure regional outage Option C is incorrect since the policy is set as Automatic, so the databases can't be failed over manually.

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

NEW QUESTION: 181

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.

A company plans to use Apache Spark analytics to analyze intrusion detection data.

What should you recommend?

- A. Azure Data Lake Storage
- B. Azure Databricks
- C. Azure HDInsight
- D. Azure Data Factory

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

Azure DataBricks does have integration with Azure Monitor. Application logs and metrics from Azure Databricks can be send to a Log Analytics workspace.

Reference:


<https://docs.microsoft.com/en-us/azure/architecture/databricks-monitoring/application-logs>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 182

You have an Azure SQL database named DB3.

You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click the Exhibit tab.)



Which Transact-SQL command should you run?

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

Answer: C (LEAVE A REPLY)

The exhibit displays Database [State] properties.

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-properties-options-page>

NEW QUESTION: 183

Case Study 3 - Contoso, Ltd 2

Overview

Contoso, Ltd. is a clothing retailer based in Seattle. The company has 2,000 retail stores across the United States and an emerging online presence.

The network contains an Active Directory forest named contoso.com. The forest is integrated with an Azure Active Directory (Azure AD) tenant named contoso.com. Contoso has an Azure subscription associated to the contoso.com Azure AD tenant.

Existing Environment

Transactional Data

Contoso has three years of customer, transaction, operational, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers.

The SQL Server instances contain data from various operations systems. The data is loaded into the instances by using SQL Server Integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time periods. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses.

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.

Streaming Twitter Data

The ecommerce department at Contoso develops an Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

Planned Changes and Requirements

Planned Changes

Contoso plans to implement the following changes:

- Load the sales transaction dataset to Azure Synapse Analytics.

- Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.

- Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

- Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.

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

- Implement a surrogate key to account for changes to the retail store addresses.

- Ensure that data storage costs and performance are predictable.

- Minimize how long it takes to remove old records.

Customer Sentiment Analytics Requirements

Contoso identifies the following requirements for customer sentiment analytics:

- Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own Azure AD credentials.

- Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.

- Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.

- Ensure that the data store supports Azure AD-based access control down to the object level.

- Minimize administrative effort to maintain the Twitter feed data records.

- Purge Twitter feed data records that are older than two years.

Data Integration Requirements

Contoso identifies the following requirements for data integration:

- Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

- Identify a process to ensure that changes to the ingestion and transformation activities can be version-controlled and developed independently by multiple data engineers.

Hotspot Question

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.

Answer Area

Table type to store retail store data:

Hash

Replicated

Round-robin

Table type to store promotional data:

Hash

Replicated

Round-robin

 Microsoft

Answer:

Answer Area  Microsoft

Table type to store retail store data:

Hash

Replicated

Round-robin

Table type to store promotional data:

Hash

Replicated

Round-robin

Explanation:

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.

Incorrect Answers:

Replicated: Replicated tables eliminate the need to transfer data across compute nodes by replicating a full copy of the data of the specified table to each compute node. The best candidates for replicated tables are tables with sizes less than 2 GB compressed and small dimension tables.

Reference:

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

NEW QUESTION: 184

SIMULATION

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.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

enable compression on a table or index

Use SQL Server Management Studio

Step 1: In Object Explorer, expand the database that contains the table that you want to compress and then expand the Tables folder.

Step 2: To compress an index, expand the table [Here: SalesLT.SalesOrderHeader table]. that contains the index that you want to compress and then expand the Indexes folder.

Step 3: In the Data Compression Wizard, on the Welcome to the Data Compression Wizard page, select Next.

Step 4: On the Select Compression Type page, select the compression type [Here: page compression] to apply to each partition in the table or index you want to compress. When finished, select Next.

Step 4a: Compression Type

Select the compression option for each partition. Isn't available when Use the same compression type for all partitions is selected. List options are None, Row, and Page[This one].

Step 4b: Skip the other options

Step 5: Right-click the table or index [Here: PK_SalesOrderHeader_SalesOrderID] to compress, point to Storage and select Manage Compression....

Step 6: In the Select an Output Option page, specify how you want to complete your compression. Select Create Script to create a SQL script based the previous pages in the wizard. Select Run immediately to create the new partitioned table after completing all remaining pages in the wizard.

Step 7: On the Review Summary page, under Review your selections, expand all available options to verify that all compression settings are correct. If everything is as expected, select Finish.

On the Compression Wizard Progress page, monitor status information about the actions of the Create Partition Wizard. Depending on the options that you selected in the wizard, the progress page might contain one or more actions. The top box displays the overall status of the wizard and the number of status, error, and warning messages that the wizard has received.)

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/data-compression/enable-compression-on-a-table-or-index?>

NEW QUESTION: 185

SIMULATION

You need to configure db1 to pause automatically after one hour of inactivity.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Azure SQL elastic pool pausing feature

An elastic pool vCore provisioned server can be paused. You can do this by changing the compute tier of the elastic pool to Serverless. In the serverless compute tier, the database is automatically paused when it is inactive for a period of time that you specify. The database is automatically resumed when the next login or other activity occurs.

Part 1: Change to the serverless compute tier

To change the compute tier of an elastic pool, you can follow these steps:

Step 1: Go to the Azure portal and sign in to your account.

Step 2: Select the Azure SQL resource that contains the elastic pool.

Step 3: In the left navigation menu, select Elastic Pools.

Step 4: Select the elastic pool that you want to change the compute tier for.

Step 5: In the Settings section, select Compute Tier.

Step 6: Select Serverless.

Step 7: Click Save.

Part 2: Configure the Auto-pause setting.

Step 1: Select db1.

Once you have changed the compute tier to serverless, the database will be paused after the idle time that you specified. You can view the idle time in the Auto-pause delay setting.

Step 2: Open database settings.

Step 3: Change the Auto-pause delay setting to 1 hour.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview>

<https://learn.microsoft.com/en-us/azure/azure-sql/database/single-database-create-quickstart>

NEW QUESTION: 186

Hotspot Question

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

VM1: 

VM2:

Answer:

Answer Area 

VM1:

VM2:

Explanation:

Box 1: Ebdsv5-series

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

Consider the following VM series based on your SQL Server workloads:

High I/O throughput-to-vCore ratio: The Mbsv3 and Mbdsbv3-series VMs offer the highest throughput-to-vCore ratio of any VM series in any cloud. The Ebdsv5-series VMs offer the second highest throughput-to-vCore ratio in Azure. If you don't know the I/O requirements for your SQL Server workload, the Ebdsv5-series VMs is the one most likely to meet your needs.

Box 2: Mv2 series

VM2 must provide the highest memory-to-vCore ratio.

The Azure Virtual Machine Mv2 series offers the highest memory-to-vCore ratio among Azure VM sizes. These VMs are specifically designed for memory-intensive workloads like in-memory databases and large-scale analytics, providing up to 12TB of memory with a substantial number of vCPUs.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices-vm-size?view=azuresql>

<https://learn.microsoft.com/en-us/azure/virtual-machines/mv2-series>

NEW QUESTION: 187

Hotspot Question

You are building a multitenant software as a service (SaaS) app named App1 that stores tenant data in an Azure SQL database.

App1 has the following requirements:

- Support a high volume of read and write requests per tenant.
- Support strict data isolation for specific tenants.

You need to recommend a sharding solution that meets the following requirements:

- Ensures that each tenant with strict data isolation can be placed in its own database

- Groups multiple tenants into a single database for cost efficiency

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

Place the tenants by using:

- Functional partitioning that separates datasets by bounded business contexts
- A list shard map and one key that maps to one shardlet
- A range shard map and contiguous keys that map to one shardlet
- Vertical partitioning that splits frequently and includes rarely used columns

Group the tenants by using:

- Functional partitioning that separates orders from databases that contain invoices
- List shardlets so that each tenant always gets a dedicated shard
- Range shardlets to group contiguous tenant IDs together
- Vertical partitioning that moves cold columns into separate tables

Answer:

Answer Area

Place the tenants by using:

- Functional partitioning that separates datasets by bounded business contexts
- A list shard map and one key that maps to one shardlet
- A range shard map and contiguous keys that map to one shardlet
- Vertical partitioning that splits frequently and includes rarely used columns

Group the tenants by using:

- Functional partitioning that separates orders from databases that contain invoices
- List shardlets so that each tenant always gets a dedicated shard
- Range shardlets to group contiguous tenant IDs together
- Vertical partitioning that moves cold columns into separate tables

Explanation:

Box 1: A list shard map and one key that maps to one shardlet.

A list shard map with a one-key-to-one-shardlet mapping is an excellent strategy for this requirement.

In Azure SQL Database Elastic Scale, this design perfectly aligns with the single-tenant (isolated) model needed for strict data compliance.

Guaranteed Isolation: Each strict-isolation tenant gets a unique tenant ID (the shard key). This key maps directly to its own dedicated database (the shardlet), ensuring zero data co-mingling.

High Performance: Dedicated databases eliminate the "noisy neighbor" effect. This allows each isolated tenant to utilize full, unshared compute resources (DTUs/vCores) to handle high read/write volumes.

Granular Scaling: You can scale resources up or down for specific high-volume tenants without affecting other users.

Easy Schema Management: Elastic Database tools allow you to run schema updates across all shards simultaneously, despite the physical separation.

Box 2: List shardlets so that each tenant always get a dedicated shard.

To support strict data isolation for specific tenants while grouping multiple other tenants into a single database for cost efficiency, you should use list shardlets (also known as list-based sharding).

The system requires a hybrid model.

Premium tenants need strict data isolation (their own dedicated database). Standard tenants need to be grouped together in shared databases to minimize costs.

List Shardlets: This strategy maps specific tenant IDs (keys) to specific databases. It allows you to explicitly assign a high-isolation tenant to its own database shardlet, while mapping a collection of standard tenant IDs to a shared database shardlet. This is the standard pattern used in Azure SQL Database Elastic Scale tools.

Incorrect:

Functional Partitioning: This separates data based on business function or subsystems (e.g., putting inventory data in one DB and billing data in another). It does not address multi-tenant isolation requirements.

Range Shardlets: This strategy groups data based on a range of keys (e.g., Tenant IDs 1-100 in Database A, 101-200 in Database B). It does not allow you to easily isolate a single, specific tenant into its own database without disrupting the continuous range logic.

Vertical Partitioning: This involves splitting a single table's columns across multiple databases (e.g., storing names in one DB and credit card info in another). It does not isolate tenants from each other.

Reference:

<https://learn.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning-strategies>

<https://www.everpuredata.com/knowledge/what-is-database-sharding.html>

NEW QUESTION: 188

You have an Azure subscription linked to an Azure Active Directory (Azure AD) tenant. The subscription contains 10 virtual machines that run Windows Server 2019 and host Microsoft SQL Server 2019 instances.

You need to ensure that you can manage the SQL Server instances by using a single user account.

What should you do first?

- A. Enable a user-assigned managed identity on each virtual machine.
- B. Deploy an Azure Active Directory Domain Services (Azure AD DS) domain and join the virtual machines to the domain.
- C. Enable a system-assigned managed identity on each virtual machine.
- D. Join the virtual machines to the Azure AD tenant.

Answer: A (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>

NEW QUESTION: 189

You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container.

Which resource provider should you enable?

- A. Microsoft.EventHub
- B. Microsoft.EventGrid
- C. Microsoft.Sql
- D. Microsoft.Automation

Answer: (SHOW ANSWER)

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger>

NEW QUESTION: 190

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

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.



Role Assignments

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



What should you use to migrate the PostgreSQL database?

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

Answer: C (LEAVE A REPLY)

Database Migration Service is a fully managed service designed to enable seamless migrations from multiple database sources to Azure data platforms with minimal downtime (online migrations).

You can use Azure Database Migration Service to migrate the databases from an on-premises PostgreSQL instance to Azure Database for PostgreSQL with minimal downtime to the application.

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

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

Downtime during the PostgreSQL database migration must be minimized.

Reference:

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

<https://docs.microsoft.com/en-us/azure/dms/tutorial-postgresql-azure-postgresql-online-portal>

NEW QUESTION: 191

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 different subnet on a virtual network. Deploy a basic Azure load balancer.
- B.** Connect each virtual machine to a different subnet on a single virtual network.
- C.** Connect each virtual machine to a single subnet on a single virtual network.
- D.** Connect each virtual machine to a single subnet on a virtual network. Deploy a standard Azure load balancer.

Answer: C (LEAVE A REPLY)

To match the on-premises experience for connecting to your availability group listener, deploy your SQL Server VMs to multiple subnets within the same virtual network. Having multiple subnets negates the need for the extra dependency on an Azure Load Balancer, or a distributed network name (DNN) to route your traffic to your listener. ref: <https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-overview?view=azuresql-vm#connectivity>

NEW QUESTION: 192

SIMULATION

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:

How to increase the maximum storage size of the query store for an azure database using Azure portal?

How to increase the maximum storage size of the query store for an azure database using SQL Server Management Studio?

To increase the size of Azure Query Store, you can use Transact-SQL within either SSMS (SQL Server Management Studio or the Azure portal. Specifically, you'll need to modify the MAX_STORAGE_SIZE_MB setting for the database. No GUI options exist for direct modification within the Azure portal.

Step 1: Connect to your database:

Use SSMS or the Azure portal to connect to the database where you want to increase the Query Store size. [Connect to db1] Step 2: . Execute the ALTER DATABASE statement:

Run the following Transact-SQL query (replacing YourDB with your database name and 1024 with the desired storage size in MB):

```
ALTER DATABASE [YourDB] SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 1024);
```

*-> Use the following query: ALTER DATABASE db1 SET QUERY_STORE

```
(MAX_STORAGE_SIZE_MB = 4000);
```

Step 3: Apply changes: The Query Store will automatically adjust to the new size.

Reference:

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

NEW QUESTION: 193

You have an Azure subscription.

You plan to deploy an instance of SQL Server on Azure Virtual Machines by using an Azure Marketplace image.

You need to register the SQL Server IaaS Agent extension (SqlIaaSExtension). The solution must meet the following requirements:

- Install critical updates for SQL Server automatically.
- Minimize performance impact on the virtual machine.

Which management mode should you select?

- A.** full
- B.** lightweight
- C.** NoAgent

Answer: (SHOW ANSWER)

Full mode installs the SQL IaaS Agent to the VM to deliver full functionality. Use it for managing a SQL Server VM with a single instance. Full mode installs two Windows services that have a minimal impact to

memory and CPU.

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extension-automate-management?view=azuresql&tabs=azure-powershell#management-modes>

NEW QUESTION: 194

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: A,E (LEAVE A REPLY)

A: Money takes less space compared to VARCHAR(20)

E: Date takes less space compared to Datetime.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/data-types/data-types-transact-sql>

NEW QUESTION: 195

Hotspot Question

You have a new Azure SQL managed instance.

You plan to create a SQL Server Agent job named SalesTotal.

You need to ensure that SalesTotal can send out notifications.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
EXECUTE msdb.dbo.sysmail_add_account_sp
```

```
    @account_name = 'SQL Agent Account',  
    @description = 'Mail account',  
    @email_address = 'salesdirector@contoso.com',  
    @display_name = 'SQL Agent Account',  
    @mailserver_name = 'smtp1.contoso.com',
```

```
EXECUTE msdb.dbo.sysmail_add_profile_sp
```

```
    @profile_name =   
        'AzureManagedInstance_dbmail_profile',  
        'AzureManagedInstanceDefault_profile',  
        'AzureManagedInstanceSystemEmail_Profile',
```

```
    @description = 'E-mail profile.';
```

```
EXEC msdb.dbo.sp_add_job @job_name=N'SalesTotals',
```

```
    @enabled=1,
```

```
    @notify_level_eventlog=0,
```

```
    @notify_level_netsend=2,  
    @notify_level_email=2,  
    @notify_level_pager=2,
```

```
    @delete_level=0,
```

```
    @description=N'Find Sales Totals',
```

```
    @category_name=N'[Uncategorized (Local)]',
```

```
    @owner_login_name=N'sa'
```

Answer:

Answer Area



```
EXECUTE msdb.dbo.sysmail_add_account_sp
```

```
@account_name = 'SQL Agent Account',  
@description = 'Mail account',  
@email_address = 'salesdirector@contoso.com',  
@display_name = 'SQL Agent Account',  
@mailserver_name = 'smtp1.contoso.com',
```

```
EXECUTE msdb.dbo.sysmail_add_profile_sp
```

```
@profile_name =
```

```
'AzureManagedInstance_dbmail_profile',  
'AzureManagedInstanceDefault_profile',  
'AzureManagedInstanceSystemEmail_Profile',
```

```
@description = 'E-mail profile';
```

```
EXEC msdb.dbo.sp_add_job @job_name=N'SalesTotals',
```

```
@enabled=1,
```

```
@notify_level_eventlog=0,
```

```
@notify_level_netsend=2,
```

```
@notify_level_email=2,
```

```
@notify_level_pager=2,
```

```
@delete_level=0,
```

```
@description=N'Find Sales Totals',
```

```
@category_name=N'[Uncategorized (Local)]',
```

```
@owner_login_name=N'sa'
```

Explanation:

Box 1: 'AzureManagedInstance_dbmail_profile',

Job notifications

SQL Agent jobs enable you to get notifications when the job finishes successfully or fails. You can receive notifications via email.

As an example exercise, set up an email account for sending the email notifications. Assign the account to the email profile called AzureManagedInstance_dbmail_profile. To send e-mail using SQL Agent jobs in SQL Managed Instance, there should be a profile that must be called AzureManagedInstance_dbmail_profile. Otherwise, SQL Managed Instance is unable to send emails via SQL Agent.

Box 2: notify_level_email

You can modify any SQL Agent job and assign operators that are notified via email if the job completes, fails, or succeeds. Modify the job by using SSMS or the following T-SQL script:

```
EXEC msdb.dbo.sp_update_job @job_name=N'Load data using SSIS',
```

```
@notify_level_email=3, -- Options are: 1 on succeed, 2 on failure, 3 on complete
```

```
@notify_email_operator_name=N'AzureSQLTeam';
```

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance>

NEW QUESTION: 196

You have an Azure SQL database named DB1.

You need to identify long-running queries against DB1 by using an Extended Events session.

Which event should you use to create the session?

- A. sqlserver.batch_completed
- B. sqlserver.sql_statement_completed
- C. sqlserver.sp_statement_completed
- D. sqlserver.rpc_completed

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

To identify long-running queries in an Azure SQL Database using Extended Events, the primary event to use is sql_statement_completed.

Reference:

https://richardswinbank.net/admin/identify_long-running_queries

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 197

You have an Azure SQL database named DB1.

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

During a database audit, you discover that none of the fields contain more than 150 characters.

You need to ensure that you can apply page compression to Table1.

What should you do?

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

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

We reduce the max length of the column from 400 to 200.

Incorrect:

Not A: Sparse column is useful when there are many null columns.

The SQL Server Database Engine uses the SPARSE keyword in a column definition to optimize the storage of values in that column. Therefore, when the column value is NULL for any row in the table, the values require no storage.

Not B, Not C: SQL Server 2005 got around the limitation of 8KB storage size and provided a workaround with varchar(max). It is a non-Unicode large variable-length character data type and can store a maximum of $2^{31}-1$ bytes (2 GB) of non-Unicode characters.

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/>

<https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/>

NEW QUESTION: 198

You are modifying an existing disaster recovery solution for an Azure SQL managed instance that contains a failover group named FG1.

You need to ensure the maximum in-transit time for FG1 when an automatic failover occurs.

What should you configure?

- A. a grace period
- B. a secondary managed instance
- C. a failover policy
- D. an availability group

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

NEW QUESTION: 199

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

<https://learn.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/automated-patching?view=azuresql> Choose only one option to automatically update your VM as overlapping tools may lead to failed updates.

NEW QUESTION: 200

SIMULATION

You need to ensure that db1 supports a minimum of one vCore, a maximum of four vCores, and a database of up to 45 GB.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

vCore purchasing model - Azure SQL Database

Select hardware configuration

You can select hardware configuration for a database or elastic pool in SQL Database at the time of creation. You can also change hardware configuration of an existing database or elastic pool.

To select a hardware configuration when creating a SQL Database or pool Step 1: Start Azure portal Step 2: Go the Azure SQL page, select db1.

Step 3: On the Overview page, select the Pricing tier link:

Microsoft Azure

Search resources, services, and docs (G+/)

Home > contoso (contoso-logicalserver/contoso) SQL database

Search

Copy Restore Export Set server firewall Delete Connect with... Feedback

This database was just created. Do you need any help getting started?

Essentials

Resource group (move) : contoso-rg

Status : Online

Location : West US 3

Subscription (move) : contoso-dev

Subscription ID : abcdef01-2345-6789-0abc-def012345678

Tags (edit) : [Click here to add tags](#)

Server name : contoso-logicalserver.database.windows.net

Elastic pool : No elastic pool

Connection strings : [Show database connection strings](#)

Pricing tier : General Purpose: Standard-series (Gen5)_2 vCores

Earliest restore point : 2023-03-20 19:21 UTC

JSON View

Step 4: Select the Change configuration link:

Microsoft Azure

Search resources, services, and docs (G+/)

Home > SQL databases > Create SQL Database >

Configure

Feedback

Service and compute tier

Select from the available tiers based on the needs of your workload. The vCore model provides a wide range of configuration controls and offers Hyperscale and Serverless to automatically scale your database based on your workload needs. Alternately, the DTU model provides set price/performance packages to choose from for easy configuration. [Learn more](#)

Service tier: General Purpose (Most budget friendly, Serverless compute)

Compare service tiers

Compute tier:

- ☒ **Provisioned** - Compute resources are pre-allocated. Billed per hour based on vCores configured.
- ☐ **Serverless** - Compute resources are auto-scaled. Billed per second based on vCores used.

Compute Hardware

Select the hardware configuration based on your workload requirements. Availability of compute optimized, memory optimized, and confidential computing hardware depends on the region, service tier, and compute tier.

Hardware Configuration: Standard-series (Gen5)

Cost summary

General Purpose (GP_Gen5_2)

Cost per vCore (in USD)

vCores selected

Cost per GB (in USD)

Max storage selected (in GB)

ESTIMATED COST / MONTH

up to 80 vCores, up to 415.23 GB memory
[Change configuration](#)

Save money
 Already have a SQL Server License? Save with a license you already own with Azure Hybrid Benefit. Actual savings may vary based on region and performance tier. [Learn more](#)

☐ Yes ☒ No

vCores [Compare vCore options](#)

2

Data max size (GB) ⓘ

32

9.6 GB LOG SPACE ALLOCATED

Apply

Step 5: Select the desired hardware configuration:

Specify:

Max vCores: 4

Max storage: 45 GB

Home > SQL databases > Create SQL Database > Configure >

SQL hardware configuration

SQL database

Available hardware configurations

Based on your workload requirements, select from available hardware configurations listed below.

Configuration	Description	Max vCores	Max memory	Max storage
Standard-series (Gen5)	Balanced memory and compute	80	415.23 GB	4 TB
M-Series	Memory optimized	--	--	--
FSv2-series	Compute optimized	--	--	--
DC-series	Enables confidential computing	--	--	--

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/service-tiers-sql-database-vcore>

NEW QUESTION: 201

You have an on-premises Microsoft SQL Server 2019 server that hosts a database named DB1.

You have an Azure subscription that contains an Azure SQL managed instance named SQLMI1 and a virtual network named VNET1. SQLMI1 resides on VNET1. The on-premises network connects to VNET1 by using an ExpressRoute connection.

You plan to migrate DB1 to SQLMI1 by using Azure Database Migration Service.

You need to configure VNET1 to support the migration.

What should you do?

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

Answer: ([SHOW ANSWER](#))

During virtual network setup, if you use ExpressRoute with network peering to Microsoft, add the following service endpoints to the subnet in which the service will be provisioned:

- * Target database endpoint (for example, SQL endpoint, Cosmos DB endpoint, and so on)
- * Storage endpoint
- * Service bus endpoint

This configuration is necessary because Azure Database Migration Service lacks internet connectivity.

Reference:

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

NEW QUESTION: 202

You have an Azure subscription. The subscription contains three instances of SQL Server on Azure Virtual Machines that run Windows Server. The subscription contains an Azure load balancer named LB1.

You implement an Always On availability group named AO1. You deploy a listener named AG1 and link AO1 to LB1.

You need to configure a load balancing rule for LB1. The solution must ensure that traffic can flow to AG1.

What should you use as the backend port?

- A. 1433
- B. 1434
- C. 59999
- D. 58888

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

The backend port for the Azure Load Balancer rule must be 1433 (or your configured SQL Server port), while the health probe must monitor port 59999 to successfully route traffic to the active Always On availability group replica.

Reference:

<https://multicloudsagar.com/file-server-cluster-in-azure/>

NEW QUESTION: 203

You have four Azure virtual machines hosted in different availability zones in the same Azure region.

You plan to deploy a Microsoft SQL Server Always On failover cluster instance (FCI) to the four virtual machines.

What is the most resilient quorum option for the deployment?

- A. a cloud witness
- B. Node and Disk Majority
- C. Node and File Share Majority
- D. Node Majority

Answer: C (LEAVE A REPLY)

A quorum mode is configured at the WSFC cluster level that dictates the methodology used for quorum voting. The Failover Cluster Manager utility will recommend a quorum mode based on the number of nodes in the cluster.

When using an asymmetric storage configuration for Always On availability groups, you should generally use the Node Majority quorum mode when you have an odd number of voting nodes, or the Node and File Share Majority quorum mode when you have an even number of voting nodes.

Reference:

<https://learn.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/wsfc-quorum-modes-and-voting-configuration-sql-server>

NEW QUESTION: 204

Your company wants to create an Azure SQL Database. The database will contain a table that would have Employee information. There are columns which contain sensitive information. You have to enable Always Encryption for the data stored in the Employee table.


Which of the following could be used to store the encryption keys required for enabling Always Encrypted for the table?

- A. Windows certificate store
- B. Private Encryption keys
- C. Public Encryption keys
- D. Azure Key vault

Answer: A,D (LEAVE A REPLY)

The Encryption keys can only be stored in either Azure Key vault or the windows certificate store.

From the implementation shown below, you can see that these are the sources you can use for storing the encryption keys.



<https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=azure-powershell>

NEW QUESTION: 205

SIMULATION

You need to protect all the databases on sql12345678 from SQL injection attacks.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Configure Advanced Threat Protection for Azure SQL Database

Advanced Threat Protection for Azure SQL Database detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Advanced Threat Protection can identify Potential *SQL injection*, Access from unusual location or data center, Access from unfamiliar principal or potentially harmful application, and Brute force SQL credentials.

Set up Advanced Threat Protection in the Azure portal

Step 1: Sign into the Azure portal.

Step 2: Navigate to the configuration page of the server you want to protect. [Here sql12345678] In the security settings, select Microsoft Defender for Cloud.

Step 3: On the Microsoft Defender for Cloud configuration page:

Step 3a: If Microsoft Defender for SQL hasn't yet been enabled, select Enable Microsoft Defender for SQL.

Step 3b: Select Configure.

The screenshot shows the Microsoft Azure portal interface. At the top, there's a blue header with the Microsoft Azure logo, a search bar, and user information (chrisqpublic@contoso.c...). Below the header, the breadcrumb trail shows 'Home > chrisqpublictest'. The main content area is titled 'chrisqpublictest | Microsoft Defender for Cloud' with a sub-label 'SQL server'. A left-hand navigation pane lists various security and performance options, with 'Microsoft Defender for Cloud' currently selected. The main content area displays a summary of security metrics: 5 Recommendations (indicated by a red exclamation mark), 0 Security alerts (indicated by a blue shield), and 5 Findings (indicated by a green shield). A prominent red-bordered box highlights the text 'Microsoft Defender for SQL: Disabled (Configure)'. Below this, a section titled 'Microsoft Defender for SQL' provides information about the service, including a 30-day trial offer. A blue button with a red border, labeled 'Enable Microsoft Defender for SQL', is prominently displayed at the bottom of the section.

Step 3c: Under ADVANCED THREAT PROTECTION SETTINGS, select Add your contact details to the subscription's email settings in Defender for Cloud.

Microsoft Azure Search resources, services, and docs (G+)

Home > chrisqpublictest >

Server settings

chrisqpublictest

Save Discard Feedback

VULNERABILITY ASSESSMENT SETTINGS

Subscription
Contoso Team
[Select Subscription](#)

Storage account
[Select Storage account](#)

Periodic recurring scans
☐ ON ☒ OFF

Send scan reports to ⓘ

☒ Also send email notification to admins and subscription owners ⓘ

ADVANCED THREAT PROTECTION SETTINGS

Advanced Threat Protection for SQL alerts emails are sent by Defender for Cloud.
[Add your contact details to the subscription's email settings in Defender for Cloud.](#) ⓘ

ⓘ Enable Auditing for better threats investigation experience

Step 3d: Provide the list of emails to receive notifications upon detection of anomalous database activities in the Additional email addresses (separated by commas) text box. [Skip] Step 3e: Optionally customize the severity of alerts that will trigger notifications to be sent under Notification types. [Skip] Step 3f: Select Save.

Microsoft Azure Search resources, services, and docs (G+)

Home > chrisqpublictest > Server settings >

Email notifications

Contoso Team

Save

Email recipients

Select who'll get the email notifications from Defender for Cloud for the Contoso Team subscription.

All users with the following roles

Owner

Additional email addresses (separated by commas)

One or more email addresses separated by commas

Notification types

Use the settings below to select the type of email notifications to be sent by Defender for Cloud.

☒ Notify about alerts with the following severity (or higher):

High

i You'll receive a maximum of one email per 6 hours for high-severity alerts, one email per 12 hours for medium-severity alerts, and one email per 24 hours for low-severity alerts. [Learn more >](#)

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/threat-detection-configure>

NEW QUESTION: 206

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 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](#))

If the number of logical processors is greater than eight (8), use eight data files. If contention continues, increase the number of data files by multiples of four (4) up to the number of logical processors until the contention is reduced to acceptable levels. Alternatively, make changes to the workload or code.

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

NEW QUESTION: 207

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](#))

You can use Azure Site Recovery to test failover.

You run a test failover to validate your replication and disaster recovery strategy, without any data loss or downtime. A test failover doesn't impact ongoing replication, or your production environment. You can run a test failover on a specific virtual machine (VM), or on a recovery plan containing multiple VMs.

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure>

NEW QUESTION: 208

Hotspot Question

You have an Azure subscription.

You plan to migrate 10 on-premises Microsoft SQL Server instances to Azure.

You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs.

The solution must minimize administrative effort.

Which SQL deployment options should you select as the master server (MSX) and the target server (TSX)? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

MSX:

▼

SQL database

SQL managed instances

SQL virtual machines

TSX:

▼

SQL database

SQL managed instances

SQL virtual machines

Answer:

Answer Area

MSX:

▼

SQL database

SQL managed instances

SQL virtual machines

TSX:

▼

SQL database

SQL managed instances

SQL virtual machines

NEW QUESTION: 209

Hotspot Question
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:

Always On
Contained
Distributed

Minimum Microsoft SQL Server version:

2017
2019
2022

Answer:

Answer Area

Type of availability group:

Always On
Contained
Distributed

Minimum Microsoft SQL Server version:

2017
2019
2022

NEW QUESTION: 210

You have an Azure subscription that contains an Azure SQL database named DB1.
You configure a new data sync group as shown in the exhibit. (Click the Exhibit tab.)

Create Data Sync Group ...

Create a sync group on the hub database. Member databases and other additional settings can be configured after create in the sync group details blade. [Learn more](#)

Sync Group Name *

Sync1 ✓

Sync Metadata Database *

db1 (Server: server1) ✓

Automatic Sync

Conflict Resolution *

Hub win ✓

Hub Database Username *

db1-admin ✓

Hub Database Password *

..... ✓

☒ Use private link

i A private connection would be created and managed by Microsoft.
[Learn more](#)

i Manual approval of the private endpoint connection is required to complete sync group creation. After sync group creation has started, go to the private endpoint page to approve the connection.

You need to minimize how long it takes to perform the first sync.

What should you do?

- A. From Private endpoint connections, approve the private endpoint connection.
- B. From Private endpoint connections, create a new private endpoint connection.
- C. From Database Sync Group, select Sync.
- D. From Create Data Sync Group, change the sync frequency to 60 seconds.

Answer: (SHOW ANSWER)

Private endpoint connections for Azure SQL Database in a data sync group need to be manually approved. This approval process ensures that the connection from the data sync service to your database is secure and authorized. You'll need to approve the service-managed private endpoint connection within the "Private endpoint connections" section in the Azure portal or through PowerShell during the sync group configuration.

Reference:

https://learn.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-sql-server-configure

NEW QUESTION: 211

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: A,E (LEAVE A REPLY)

Always Encrypted isn't supported for the columns using one of the following data types: xml, timestamp, rowversion, image, ntext, text, sql_variant, hierarchyid, geography, geometry, alias, user-defined types.

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 212

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 the on-premises networks shown in the following table.

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 db1 and db2 users', '131.107.10.0',  
'131.107.10.255'
```

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 213

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](#))

B: The auto-failover groups feature allows you to manage the replication and failover of some or all databases on a logical server to another region.

C: The auto-failover group must be configured on the primary server and will connect it to the secondary server in a different Azure region.



Incorrect:

Not A: Geo-replication failovers require client applications to change their connection strings.

Reference:

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

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

NEW QUESTION: 214

You have 50 Azure SQL databases.

You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.

What should you do?

- A. Create a diagnostic setting for the activity log that has the Security log enabled.
- B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
- C. Create an alert rule that uses a Metric signal type.
- D. Create an alert rule that uses an Activity Log signal type.

Answer: D (LEAVE A REPLY)

Activity log events - An alert can trigger on every event, or, only when a certain number of events occur.

Incorrect Answers:

C: Metric values - The alert triggers when the value of a specified metric crosses a threshold you assign in either direction. That is, it triggers both when the condition is first met and then afterwards when that

condition is no longer being met.
Reference:
<https://docs.microsoft.com/en-us/azure/azure-sql/database/alerts-insights-configure-portal>

NEW QUESTION: 215

Drag and Drop Question

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change. You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00 in the morning, and transaction log backups every five minutes. The database fails on Wednesday at 10:00 in the morning. Which three backups should you restore in sequence? To answer, move the appropriate backups from the list of backups to the answer area and arrange them in the correct order.

Actions

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups

full backup

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Wednesday log backups

Wednesday differential backup

Answer Area

>

<

↑

↓



Answer:

Actions

Monday, Tuesday, and then Wednesday differential backups

Wednesday, Tuesday, and then Monday log backups



Microsoft

Monday, Tuesday, and then Wednesday log backups

Wednesday, Tuesday, and then Monday differential backups

Answer Area

full backup

Wednesday differential backup

Wednesday log backups

>

<

↑

↓

Explanation:

A differential backup is based on the most recent, previous full data backup.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/differential-backups-sql-server?view=sql-server-ver15>

NEW QUESTION: 216

You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2. SQLMI1 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.

What should you use to perform the restore?

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

Answer: C (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/point-in-time-restore?view=azuresql-mi&tabs=azure-portal> You can use point-in-time restore to restore a database in these scenarios:

From an existing database

From a deleted database

To the same managed instance or to a different managed instance

To a managed instance in the same subscription or to a managed instance in a different subscription

NEW QUESTION: 217

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

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

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

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

Solution: You restore Database1 from Server1 to the Server2 by using the RESTORE Transact- SQL command and the REPLACE option.

Does this meet the goal?

A. Yes

B. No

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

You can't overwrite an existing Azure SQL Database database during restore.

Correct 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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups>

NEW QUESTION: 218

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: A,C,E ([LEAVE A REPLY](#))

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

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

Reference:

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

NEW QUESTION: 219

Case Study 2 - Contoso, Ltd

Overview

General Overview

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

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure. The company has customers worldwide.

Existing Environment

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 he 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.



Role Assignments

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



You are evaluating the business goals.

Which feature should you use to provide customers with the required level of access based on their service agreement?

- A. dynamic data masking
- B. Conditional Access in Azure
- C. service principals
- D. row-level security (RLS)

Answer: (SHOW ANSWER)

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

Row-level security (RLS) is a feature added as of SQL Server 2016. Instead of encrypting or decrypting a database's table data, it restricts and filters a table's row-level data in accordance with security policies defined by the user. This enables the database engine to limit the number of exposed data rows.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15>

NEW QUESTION: 220

Drag and Drop Question

You have an Azure SQL Database named DB1.

You plan to create a table named Sales that will contain a column named Salesdate.

You need to implement table partitioning based on the Salesdate column. The solution must ensure that the lowest boundary value is included in each partition.

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

Actions	Answer Area
Create a table.	1
Create a partition scheme.	2
Create a file group.	3
Create a partition function that uses the AS RANGE LEFT FOR VALUES clause.	
Create a partition function that uses the AS RANGE RIGHT FOR VALUES clause.	

Answer:

Actions	Answer Area
Create a table.	1 Create a partition function that uses the AS RANGE RIGHT FOR VALUES clause.
Create a partition scheme.	2 Create a partition scheme.
Create a file group.	3 Create a table.
Create a partition function that uses the AS RANGE LEFT FOR VALUES clause.	
Create a partition function that uses the AS RANGE RIGHT FOR VALUES clause.	

Explanation:

Step 1: Create a partition function that uses the AS RANGE RIGHT FOR VALUES clause.

To partition an Azure SQL Database table by a sale date while including the lowest boundary value in each partition, you must use a RANGE RIGHT partition function.

Define the boundaries for your partitions using the RANGE RIGHT specification. This ensures that the boundary date itself acts as the inclusive lowest value of that partition (e.g., a boundary of 2026-01-01 starts a partition covering that day and onward).

```
CREATE PARTITION FUNCTION pf_SaleDate (DATE)
```

```
AS RANGE RIGHT
```

```
FOR VALUES ('2026-01-01', '2026-02-01', '2026-03-01');
```

Step 2: Create a partition scheme

```
Create Partition Scheme
```

Map the partition function to the storage destinations. In Azure SQL Database, all partitions must be mapped to the [PRIMARY] filegroup because custom filegroups are not supported.

```
CREATE PARTITION SCHEME ps_SaleDate
```

```
AS PARTITION pf_SaleDate
```

```
ALL TO ([PRIMARY]);
```

Step 3: Create a table

```
Create Table on Partition Scheme
```

Create the table and assign the partition scheme to the sale date column instead of a traditional filegroup specification.

```
CREATE TABLE Sales (
```

```
SaleID INT IDENTITY(1,1),
```

```
SaleDate DATE NOT NULL,  
Amount DECIMAL(18,2),  
CustomerID INT,  
-- The partitioning column must be part of the primary key if enforced  
CONSTRAINT PK_Sales PRIMARY KEY CLUSTERED (SaleID, SaleDate)  
)  
ON ps_SaleDate (SaleDate);  
Reference:
```

<https://techcommunity.microsoft.com/blog/modernizationbestpracticesblog/part-2---azure-sql-db-hyperscale-table-partitioning---best-practices--recommenda/3721453>

NEW QUESTION: 221

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

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

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

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

Solution: You run the Remove-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabase PowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

A. Yes

B. No

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

Correct 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.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/recovery-using-backups>

NEW QUESTION: 222

You have an Azure subscription that contains a SQL Server on Linux Azure Virtual Machines instance named VM1. VM1 contains a database named DB1.

You need to create operational performance baselines for DB1 and VM1.

What should you use?

A. InfluxDB, Collectd, and Grafana

B. Azure Monitor Log Analytics

C. Azure Event Hubs

D. top, grep, and sed

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

Azure Monitor Logs provides monitoring, alerting, and alert remediation capabilities across cloud and on-premises assets. The Log Analytics virtual machine extension for Linux is published and supported by Microsoft. The extension installs the Log Analytics agent on Azure virtual machines, and enrolls virtual machines into an existing Log Analytics workspace.

NEW QUESTION: 223

Hotspot Question

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

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:

Answer Area

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.

Explanation:

Box 1: Export the data

For databases that don't qualify for reverse migration [See Limitations for reverse migration below], the only way to migrate from Hyperscale to a non-Hyperscale service tier is to export/import using a bacpac file or other data movement technologies (Bulk Copy, Azure Data Factory, Azure Databricks, SSIS, etc.) Box 2: Set the Service tier to General Purpose.

Reverse migrate a database from Hyperscale

You can migrate an existing Hyperscale database in Azure SQL Database to the General Purpose service tier using the Azure portal, the Azure CLI, PowerShell, or Transact-SQL.

Reverse migration to the General Purpose service tier allows customers who have recently converted an existing database in Azure SQL Database to Hyperscale to move back in an emergency, should Hyperscale not meet their needs. While reverse migration is initiated by a service tier change, it's essentially a size-of-data move between different architectures.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/reverse-migrate-from-hyperscale>

NEW QUESTION: 224

You have an Azure SQL managed instance.

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

What should you do?

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

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

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15>

NEW QUESTION: 225

You have an Azure subscription.

You need to deploy an Azure SQL database by using a BACPAC file.

Which command should you run?

- A.** az sql db create
- B.** az sql db op
- C.** az sql db import
- D.** az sql db copy

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

Import a BACPAC file into a database in SQL Database using the Azure CLI Example:

```
echo "Importing sample database from $container to $database..."
```

```
az sql db import --admin-password $password --admin-user $login --storage-key $key --storage- key-type StorageAccessKey --storage-uri
```

```
https://$storage.blob.core.windows.net/$container/$bacpac --name $database --resource-group
```

```
$resourceGroup --server $server
```

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/scripts/import-from-bacpac-cli>

NEW QUESTION: 226

Hotspot Question

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.

How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Encryption key: ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type: ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Answer:

Answer Area Microsoft

Encryption key: ▼

- Column encryption key
- Database encryption key
- Service master key

Encryption type: ▼

- Deterministic
- Randomized
- Transparent Data Encryption (TDE)

Explanation:

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Incorrect Answers:

TDE encrypts the storage of an entire database by using a symmetric key called the Database Encryption Key (DEK).

Box 2: Deterministic

Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to the Database Engine (SQL Database or SQL Server).

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption. Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

Incorrect Answers:

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.

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. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (**503** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 227

Hotspot Question

You have an on-premises server named Server1 that has Microsoft SQL Server 2022 installed.

You need to define an extended event to track user connections to Server1. The solution must ensure that all connection attempts are returned.

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

```
CREATE EVENT SESSION [ExtendedEvents001]
ON SERVER
ADD EVENT
```

sqlserver.login
sqlserver.logout
sqlserver.process_login_finish
sqlserver.session_fedauth_failure

```
SET collect_options_text = (1)
ACTION(sqlserver.client_app_name, sqlserver.client_connection_id,
sqlserver.client_hostname,
sqlserver.context_info, sqlserver.server_instance_name,
sqlserver.server_principal_name,
sqlserver.username)
),
```

```
ADD EVENT
```

sqlserver.login
sqlserver.logout
sqlserver.process_login_finish
sqlserver.session_fedauth_failure

```
ACTION(sqlserver.client_app_name, sqlserver.client_connection_id,
sqlserver
client_hostname, sqlserver.context_info,
sqlserver.server_instance_name,
sqlserver.server_principal_name, sqlserver.session_id,
sqlserver.username)
```

```
)
ADD TARGET package0.event_file (
SET filename = N 'ExtendedEvents001'
```

```
);
```

Answer:

Answer Area

```
CREATE EVENT SESSION [ExtendedEvents001]
ON SERVER
ADD EVENT 

|                                   |
|-----------------------------------|
| sqlserver.login                   |
| sqlserver.logout                  |
| sqlserver.process_login_finish    |
| sqlserver.session_fedauth_failure |



SET collect_options_text = (1)
ACTION(sqlserver.client_app_name, sqlserver.client_connection_id,
sqlserver.client_hostname,
sqlserver.context_info, sqlserver.server_instance_name,
sqlserver.server_principal_name,
sqlserver.username)
),
ADD EVENT 

|                                   |
|-----------------------------------|
| sqlserver.login                   |
| sqlserver.logout                  |
| sqlserver.process_login_finish    |
| sqlserver.session_fedauth_failure |



ACTION(sqlserver.client_app_name, sqlserver.client_connection_id,
sqlserver
    client_hostname, sqlserver.context_info,
sqlserver.server_instance_name,
    sqlserver.server_principal_name, sqlserver.session_id,
sqlserver.username)

)
ADD TARGET package0.event_file (
    SET filename = N 'ExtendedEvents001'

);
```

Explanation:

Box 1: sqlserver.login

This is an abbreviated version of process_login_finish, containing only the columns required by external monitoring telemetry pipeline.

Use this for Box 1 as it has fewer parameters.

Box 2: sqlserver.process_login_finish

This event is generated when server is done processing a login (success or failure).

Reference:

<https://www.sqlskills.com/blogs/bobb/over-1000-xevents-in-sql-server-2016-ctp2-here-are-the-new-ones/>

NEW QUESTION: 228

SIMULATION

You need to configure your user account as the Microsoft Entra admin for the server named sql12345678.

To complete this task, sign in to the virtual machine. You may need to use SQL Server Management Studio and the Azure portal.

Answer:

Manage administrator accounts in Azure Active Directory B2C

You can assign a role when you create a user or invite a guest user. You can add a role, change the role, or remove a role for a user:

Solution 1:

Provision Microsoft Entra admin (SQL Managed Instance)

Your SQL Managed Instance needs permission to read Microsoft Entra ID to accomplish tasks such as authentication of users through security group membership or creation of new users. For this to work, you must grant the SQL Managed Instance permission to read Microsoft Entra ID.

You can do this using the Azure portal or PowerShell.

To grant your SQL Managed Instance read permissions to Microsoft Entra ID using the Azure portal, sign in as a Global Administrator and follow these steps:


Step 1: In the Azure portal, in the upper-right corner select your account, and then choose Switch directories to confirm which directory is your Current directory. Switch directories, if necessary.



Step 2: Choose the correct Microsoft Entra directory as the Current directory.

This step links the subscription associated with Microsoft Entra ID to the SQL Managed Instance, ensuring the Microsoft Entra tenant and SQL Managed Instance use the same subscription.

Step 3: Now, you can choose your Microsoft Entra admin for your SQL Managed Instance. For that, go to your managed instance resource in the Azure portal and select Microsoft Entra admin under Settings.




Step 4: Select the banner on top of the Microsoft Entra admin page and grant permission to the current user.




Step 5: After the operation succeeds, the following notification will show up in the top-right corner:



Step 6: On the Microsoft Entra admin page, select Set admin from the navigation bar to open the Microsoft Entra ID pane.



Step 7: On the Microsoft Entra ID pane, search for a user, check the box next to the user or group to be an administrator, and then press Select to close the pane and go back to the Microsoft Entra admin page for your managed instance. [Select yourself] Step 8: From the navigation bar of the Microsoft Entra admin page for your managed instance, select Save to confirm your Microsoft Entra administrator.



The process of changing the administrator might take several minutes. Then the new administrator appears in the Microsoft Entra admin box.

Solution 2:

Step 1: Sign in to the Azure portal with Global Administrator or Privileged Role Administrator permissions.

Step 2: (If you have access to multiple tenants, select the Settings icon in the top menu to switch to your Azure AD B2C tenant from the Directories + subscriptions menu.) Step 3: Under Azure services, select Azure AD B2C. Or use the search box to find and select Azure AD B2C.

Step 4: Under Manage, select Users.

Step 5: Select the user you want to change the roles for. [Select your user account]. Then select Assigned roles.

Step 6: Add a role assignment

Step 7: Select Add assignments, select the role to assign (for example, Application administrator), [Select Microsoft Entra administrator] and then choose Add.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-configure>

<https://learn.microsoft.com/en-us/azure/active-directory-b2c/tenant-management-manage-administrator>

NEW QUESTION: 229

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

You have an Azure subscription.

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

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

Which service tier should you use?

- A. Hyperscale
- B. General Purpose
- C. Premium
- D. Business Critical

Answer: (SHOW ANSWER)

General purpose has disk latency, Business critical puts the logs on SSD, Premium is a DTU offering which is not supported on Azure SQL MI, Hyperscale is not on Managed Instance.

NEW QUESTION: 230

Drag and Drop Question

You have an Azure virtual machine named Server1 that contains an instance of Microsoft SQL Server 2022 named SQL1.

SQL1 contains two databases named DB1 and DB2.

You need to take a snapshot backup of DB1 and DB2. The backup must NOT disrupt the backup chain and must NOT affect other databases.

Which four Transact-SQL commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Commands

```
ALTER DATABASE DB1
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```

```
BACKUP SERVER
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET
SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

```
BACKUP GROUP db1, db2
TO DISK = 'd:\temp\db.bkm'
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION
SET
SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER DATABASE DB2
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```

Answer Area

1

2

3

4



Answer:

Commands

```
ALTER DATABASE DB1  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```

```
BACKUP SERVER  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION  
SET  
  
SUSPEND_FOR_SNAPSHOT_BACKUP = ON;
```

```
BACKUP GROUP db1, db2  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

```
ALTER SERVER CONFIGURATION  
SET  
SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

```
ALTER DATABASE DB2  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```

Answer Area

1

```
ALTER SERVER CONFIGURATION  
SET  
SUSPEND_FOR_SNAPSHOT_BACKUP = ON  
(GROUP = (db1, db2), MODE = COPY_ONLY);
```

2

```
BACKUP GROUP db1, db2  
TO DISK = 'd:\temp\db.bkm'  
WITH METADATA_ONLY, FORMAT;
```

3

```
ALTER DATABASE DB1  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```

4

```
ALTER DATABASE DB2  
SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF;
```



Explanation:

Note: There is also another older very similar question with three steps, instead of four.

Unfreeze in two steps in this version.

Step 1: ALTER SERVER CONFIGURATION SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON (GROUP = (DB1, DB2), MODE = COPY_ONLY); ALTER SERVER CONFIGURATION ... GROUP = (DB1, DB2): Freezes the write I/O on multiple specified databases simultaneously. This keeps other databases on the server unaffected, which rules out BACKUP SERVER or instance-wide alternatives.

Suspend multiple user databases for snapshot backup

If multiple databases on the same underlying disk, you could suspend multiple databases with the following command.

ALTER SERVER CONFIGURATION

SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON

(GROUP = (testdb1, testdb2));

Step 2: BACKUP GROUP DB1, DB2 TO DISK = 'D:\Backups\GroupSnapshot.bkm' WITH METADATA_ONLY, FORMAT; Perform snapshots of multiple databases simultaneously with the new BACKUP GROUP and BACKUP SERVER commands. With this option, snapshots can be performed at the snapshot granularity of the underlying storage, eliminating the need to perform a snapshot of the same disk multiple times.

Step 3: ALTER DATABASE DB1

SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF

Resume write I/O on the DB1 to unfreeze the databases for normal operation.

In SQL Server 2022, the command ALTER DATABASE databasename SET

SUSPEND_FOR_SNAPSHOT_BACKUP = OFF unfreezes a database's write I/O operations and resumes normal database processing. It explicitly cancels or ends a manual database freeze state that was previously initiated for a storage-level backup.

Step 4: ALTER DATABASE DB2

SET SUSPEND_FOR_SNAPSHOT_BACKUP = OFF

Resume write I/O on the DB2 to unfreeze the databases for normal operation.

Reference:

<https://learn.microsoft.com/en-us/sql/relational-databases/backup-restore/create-a-transact-sql-snapshot-backup>

NEW QUESTION: 231

Hotspot Question

You have an Azure SQL database.

You are reviewing a slow performing query as shown in the following exhibit.



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

Answer Area

The exhibit shows [answer choice].

The [answer choice] operator in the execution plan indicates that the query would benefit from performance tuning.

Microsoft

an actual execution plan
an estimated execution plan
Live Query Statistics

Index Seek
Key Lookup
Nested Loops

Answer:

Answer Area

The exhibit shows [answer choice].

The [answer choice] operator in the execution plan indicates that the query would benefit from performance tuning.

Microsoft

an actual execution plan
an estimated execution plan
Live Query Statistics

Index Seek
Key Lookup
Nested Loops

Explanation:

Box 1: Live Query Statistics

Live Query Statistics as it a percentage of the execution.

Box 2: Key Lookup

The use of a Key Lookup operator in a query plan indicates that the query might benefit from performance tuning. For example, query performance might be improved by adding a covering index.

Reference:

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

<https://docs.microsoft.com/en-us/sql/relational-databases/showplan-logical-and-physical-operators-reference>

NEW QUESTION: 232

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

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

Which resource type should you use to create the alert?

A. Resource Groups

B. SQL Servers

C. SQL Databases

D. SQL Virtual Machines

Answer: ([SHOW ANSWER](#))

There are resource types related to application code, compute infrastructure, networking, storage + databases.

You can deploy up to 800 instances of a resource type in each resource group.

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

Reference:

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

NEW QUESTION: 233

You have an Azure SQL database.

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

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

What should you use?

A. Performance Recommendations in the Azure portal

B. Extended Events in Microsoft SQL Server Management Studio (SSMS)

C. Query Store in Microsoft SQL Server Management Studio (SSMS)

D. Query Performance Insight in the Azure portal

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

Use the Query Store Page in SQL Server Management Studio.

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

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

NEW QUESTION: 234

Hotspot Question

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 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:

Box 1: Create a replicate on the same logical server.

A High Availability (HA) replica uses the same page servers as the primary replica.

A named replica, just like an HA replica, uses the same page servers as the primary replica.

Box 2: Add an ApplicationIntent entry to the connection string.

Connecting to an HA replica.

In Hyperscale databases, the ApplicationIntent argument in the connection string used by the client dictates whether the connection is routed to the read-write primary replica or to a read-only HA replica. If ApplicationIntent is set to ReadOnly and the database doesn't have a secondary replica, connection will be routed to the primary replica and will default to the ReadWrite behavior.

Example:

-- Connection string with application intent

Server=tcp:<myserver>.database.windows.net;Database=<mydatabase>;ApplicationIntent=Read Only;User ID=<myLogin>;Password=<myPassword>;Trusted_Connection=False; Encrypt=True; Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/service-tier-hyperscale-replicas?view=azuresql&tabs=tsql>

<https://sqlserverguides.com/read-only-replica-azure-sql/>

NEW QUESTION: 235

Drag and Drop Question

You have an Azure subscription that contains an Azure SQL managed instance named Server1.

You need to enable auditing for Server1. The solution must minimize administrative effort.

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

Actions	Answer Area
On Server1, run the create server audit specification T-SQL statement.	
From the Azure portal, create an Azure Storage account	
On Server1, ceate a credential.	
On Server1, enable diagnostics.	
On Server1, run the create server audit T-SQL statement.	
From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.	

Answer:

Actions	Answer Area
On Server1, run the create server audit specification T-SQL statement.	From the Azure portal, create an Azure Storage account
	From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
	On Server1, create a credential.
On Server1, enable diagnostics.	On Server1, run the create server audit T-SQL statement.

NEW QUESTION: 236

Your on-premises network contains a server that hosts a 60-TB database named DB1. The network has a 10-Mbps internet connection.

You need to migrate DB1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

- A. Azure Migrate
- B. Azure Data Box
- C. Azure Database Migration Service
- D. Data Migration Assistant (DMA)

Answer: B (LEAVE A REPLY)

Database bigger than 40TB, better to use Azure Data Box.

<https://learn.microsoft.com/en-us/azure/databox/data-box-overview>

NEW QUESTION: 237

Hotspot Question

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.

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

(LOCATION = 'https://account1.

blob

dfs

table

.core.windows.net',

PUSHDOWN = ON

TYPE = BLOB_STORAGE

TYPE = HADOOP

```
)
```



Answer:

Answer Area

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

(LOCATION = 'https://account1.

blob

dfs

table

.core.windows.net',

PUSHDOWN = ON

TYPE = BLOB_STORAGE

TYPE = HADOOP

```
)
```



Explanation:

Box 1: dfs

For Azure Data Lake Store Gen 2 used the following syntax:

http[s] <storage_account>.dfs.core.windows.net/<container>/subfolders

Incorrect:

Not blob: blob is used for Azure Blob Storage. Syntax:

http[s] <storage_account>.blob.core.windows.net/<container>/subfolders

Box 2: TYPE = HADOOP

Syntax for CREATE EXTERNAL DATA SOURCE.

External data sources with TYPE=HADOOP are available only in dedicated SQL pools.

```
CREATE EXTERNAL DATA SOURCE <data_source_name>
```

```
WITH -
```

```
( LOCATION = '<prefix>://<path>'
```

```
[, CREDENTIAL = <database scoped credential> ]
```

```
, TYPE = HADOOP
```

```
)
```

```
[;]
```

Reference:

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

NEW QUESTION: 238

You have eight on-premises servers that have Microsoft SQL Server 2022 installed. Each server contains multiple databases.

You plan to migrate the on-premises databases to Azure.

You need to analyze the servers to identify which Azure SQL services can be used to host the databases.

Which two tools can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Data Migration Assistant (DMA)
- B. SQL Server Migration Assistant (SSMA)
- C. Azure Database Migration Service
- D. Azure Migrate
- E. Database Experimentation Assistant (DEA)

Answer: A,D (LEAVE A REPLY)

[A] The Data Migration Assistant (DMA) can be used to identify which Azure SQL services are suitable for hosting your databases when migrating from SQL Server 2022 to Azure. DMA assesses your on-premises SQL Server databases and provides recommendations on whether they are compatible with Azure SQL Database or Azure SQL Managed Instance. It identifies potential migration blockers and compatibility issues, helping you choose the appropriate Azure SQL service.

[D] Azure Migrate can be used to identify which Azure SQL services are suitable for hosting your migrated databases. It assesses on-premises SQL Server databases for migration readiness, providing recommendations on target Azure SQL service options like Azure SQL Database (single database or elastic pool), Azure SQL Managed Instance, and SQL Server on Azure VMs.

Azure Migrate also helps with performance-based sizing and cost estimations for running the databases in Azure.

Reference:

<https://learn.microsoft.com/en-us/sql/dma/dma-overview?view=sql-server-ver17>

<https://learn.microsoft.com/en-us/data-migration/sql-server/database/guide>

NEW QUESTION: 239

You are designing an anomaly detection solution for streaming data from an Azure IoT hub.

The solution must meet the following requirements:

- Send the output to an Azure Synapse.
- Identify spikes and dips in time series data.
- Minimize development and configuration effort.

Which should you include in the solution?

- A. Azure SQL Database
 - B. Azure Databricks
 - C. Azure Stream Analytics
- Answer: C (LEAVE A REPLY)**

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

NEW QUESTION: 240

Hotspot Question

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:

Premium block blobs

Premium file shares

Standard general-purpose v2

Replication:

Geo-redundant storage (GRS)

Locally-redundant storage (LRS)

Read-access geo-redundant storage (RA-GRS)



Answer:

Answer Area

Type:

Premium block blobs

Premium file shares

Standard general-purpose v2

Replication:

Geo-redundant storage (GRS)

Locally-redundant storage (LRS)

Read-access geo-redundant storage (RA-GRS)



NEW QUESTION: 241

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: B (LEAVE A REPLY)

Serverless is a compute tier for single databases in Azure SQL Database that automatically scales compute based on workload demand and bills for the amount of compute used per second. The serverless compute tier also automatically pauses databases during inactive periods when only storage is billed and automatically resumes databases when activity returns. The serverless compute tier is available in the General Purpose service tier and currently in preview in the Hyperscale service tier.

<https://learn.microsoft.com/en-us/azure/azure-sql/database/serverless-tier-overview?view=azuresql&tabs=general-purpose>

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 242
Hotspot Question
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(),

IsMemoryOptimizedElevateToSnapshotEnabled;

IsXTPSupported;

Updateability;

ALTER DATABASE CURRENT SET

COMPATIBILITY LEVEL

MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT

READ_COMMITTED_SNAPSHOT

= ON;

Answer:

Answer Area

```
SELECT DATABASEPROPERTYEX(DB_NAME(),
```

```
IsMemoryOptimizedElevateToSnapshotEnabled;  
IsXTPSupported;  
Updateability;
```

```
ALTER DATABASE CURRENT SET
```

```
COMPATIBILITY LEVEL  
MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT  
READ_COMMITTED_SNAPSHOT
```

```
= ON;
```

Explanation:

Box 1: IsXTPSupported;

Step 1: Ensure you are using a Premium or Business Critical tier database In-Memory OLTP is supported if the result from the following query is 1 (not 0):

```
SELECT DATABASEPROPERTYEX(DB_NAME(), 'IsXTPSupported');
```

XTP stands for Extreme Transaction Processing, which is an informal name of the In-Memory OLTP feature.

(Step 2: Identify objects to migrate to In-Memory OLTP)

Box 2: MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT

Step 3: Create a comparable test database

Suppose the report indicates your database has a table that would benefit from being converted to a memory-optimized table. We recommend that you first test to confirm the indication by testing.

You need a test copy of your production database. The test database should be at the same service tier level as your production database.

To ease testing, tweak your test database as follows:

Connect to the test database by using SQL Server Management Studio (SSMS).

To avoid needing the WITH (SNAPSHOT) option in queries, set the current database's MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT option, as shown in the following T-SQL statement:

```
ALTER DATABASE CURRENT SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT = ON;
```

(Step 4: Migrate tables)

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/in-memory-oltp-configure>

NEW QUESTION: 243

You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.

You need to migrate the databases hosted on SQL1 to Azure. The solution must meet the following requirements:

- The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections.
- Administrative effort must be minimized.

What should you use to host the databases?

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

Answer: (SHOW ANSWER)

Linked servers are available in SQL Server Database Engine and Azure SQL Managed Instance.
They are not enabled in Azure SQL Database singleton and elastic pools.
<https://learn.microsoft.com/en-us/sql/relational-databases/linked-servers/linked-servers-database-engine?view=sql-server-2017>

NEW QUESTION: 244

Drag and Drop Question
You have an Azure subscription that contains a SQL Server on Azure virtual Machines instance named SQLVM1.
You need to ensure that SQL Server Agent jobs on SQLVM1 can run operating system commands.
Which four T-SQL statements should you execute in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

⋮ USE master

⋮ CREATE CREDENTIAL

⋮ EXEC dbo.sp_grant_login_to_proxy

⋮ EXEC dbo.sp_grant_proxy_to_subsystem

⋮ USE msdb

⋮ EXEC dbo.sp_add_proxy

Answer Area

Answer:

Statements

USE master

EXEC dbo.sp_grant_login_to_proxy

Answer Area

USE msdb

CREATE CREDENTIAL

EXEC dbo.sp_add_proxy

EXEC dbo.sp_grant_proxy_to_subsystem

Valid DP-300 Dumps shared by Actual4test.com for Helping Passing DP-300 Exam! Actual4test.com now offer the **newest DP-300 exam dumps**, the Actual4test.com DP-300 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-300 dumps with Test Engine here: https://www.actual4test.com/DP-300_examcollection.html (503 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)