Number: DP-300
Passing Score: 800
Time Limit: 120
File Version: 42.0

Exam Code: DP-300
Exam Name: Administering Relational Databases on Microsoft Azure

01 - Implement a Secure Environment

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

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.

QUESTION 1
DRAG DROP
You create all of the tables and views for ResearchDB1.
You need to implement security for ResearchDB1. 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.

Select and Place:

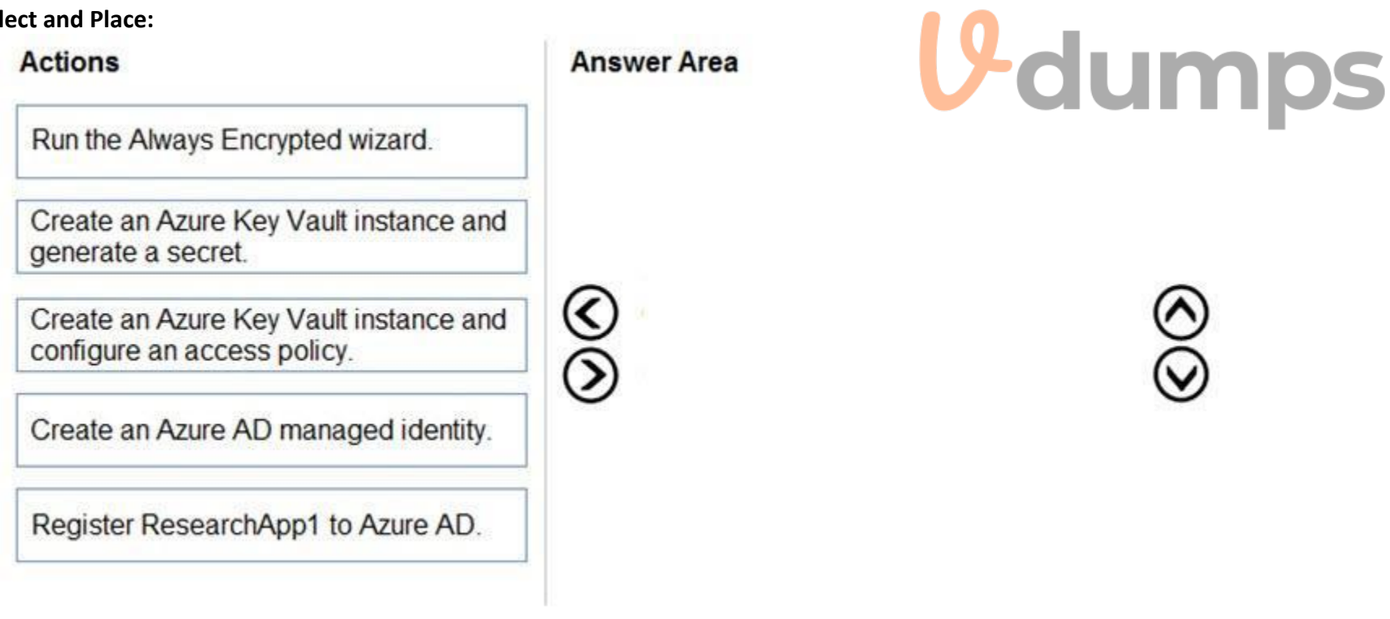

Correct Answer:

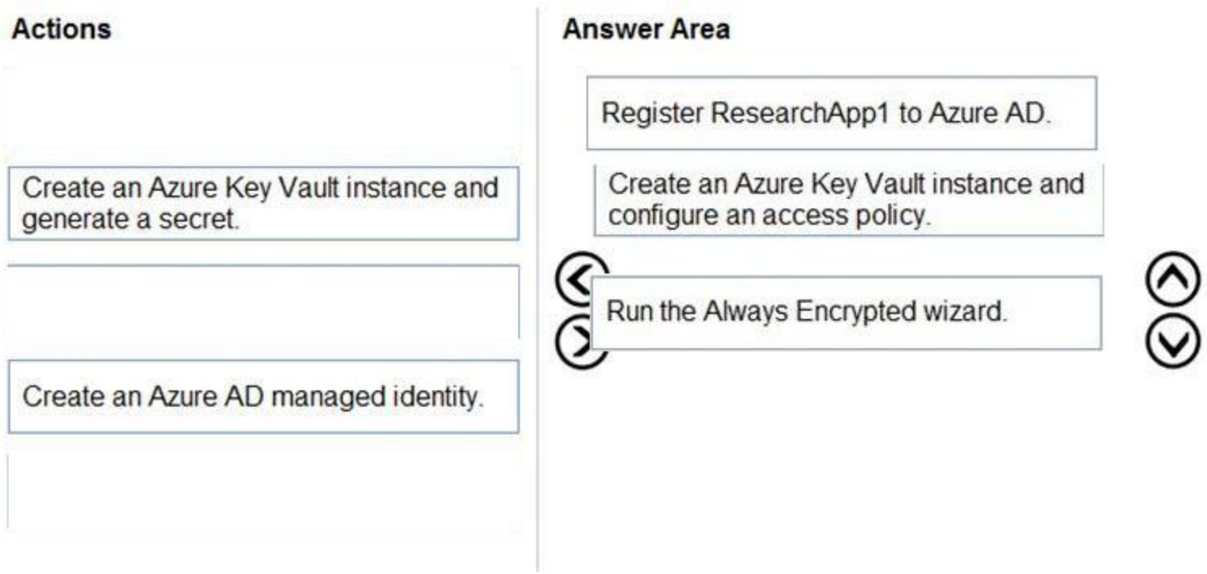

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=azure-powershell

QUESTION 2
DRAG DROP
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.

Select and Place:

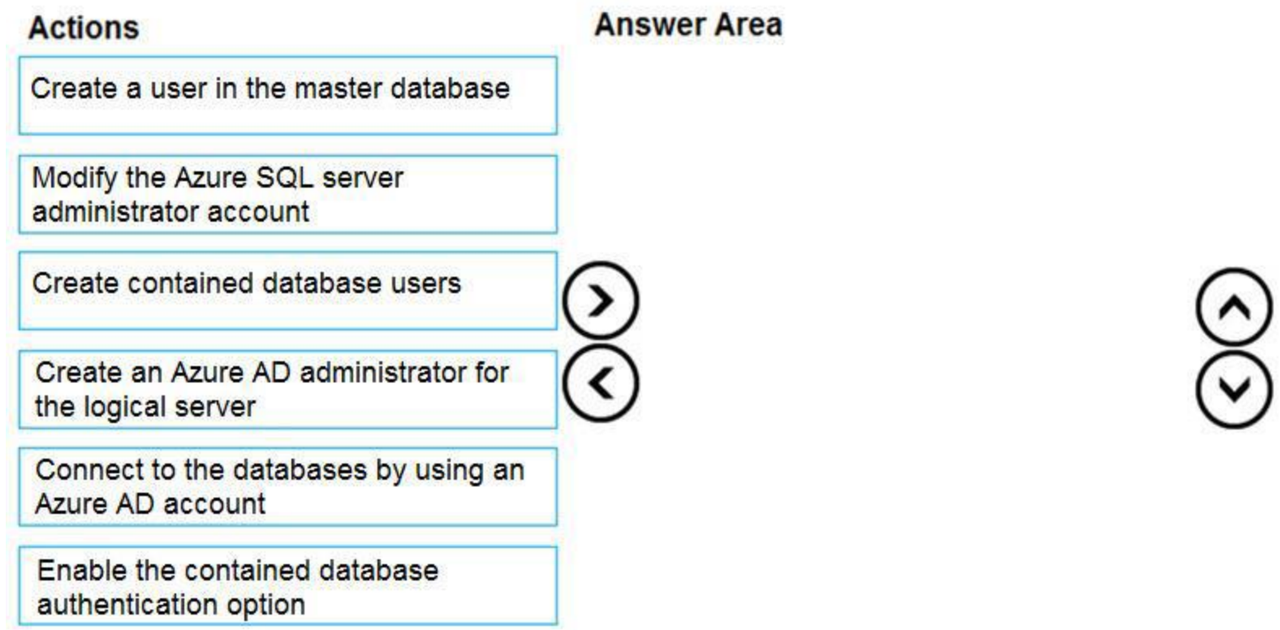

Correct Answer:

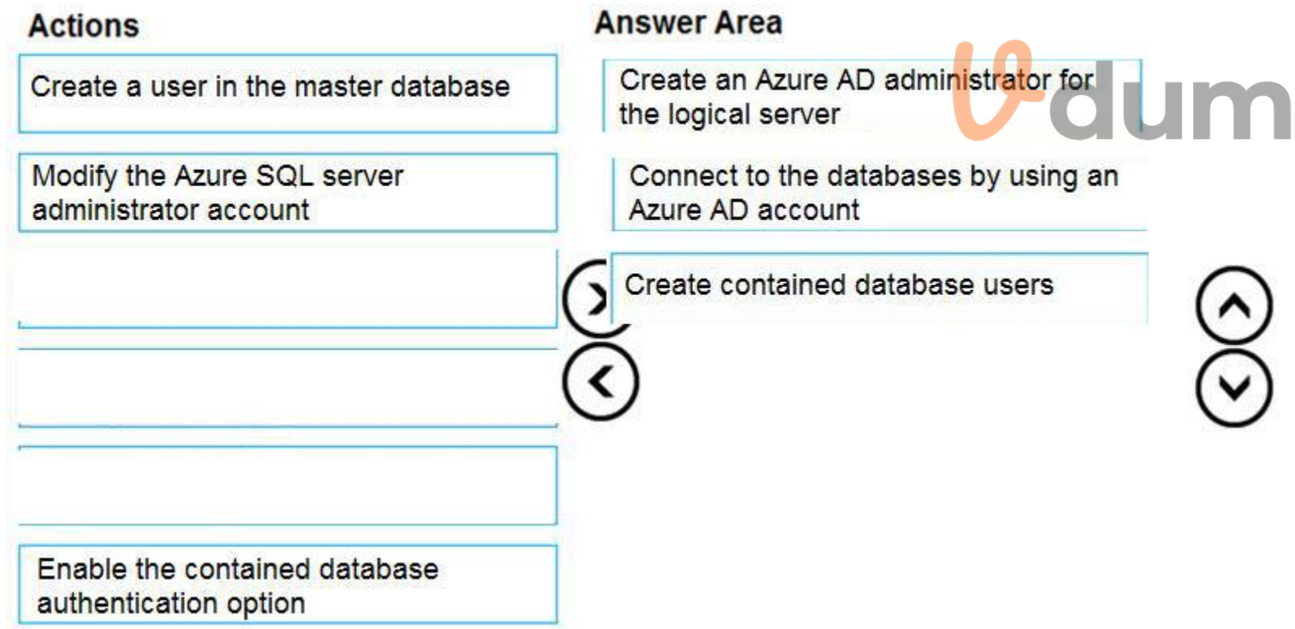

Section:

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. Connect to the databases using an Azure AD account (the Administrator account that was configured in the previous step). (Step 2)

5. Create contained database users in your database mapped to Azure AD identities. (Step 3)
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-configure?tabs=azure-powershell

02 - Implement a Secure Environment
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
Existing Environment
Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.
Active Directory
Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.
Database Environment
Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

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

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.
The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.
Current Business Model
Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers. The customers receive the files by using FTP.
Requirements
Planned Changes
Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.
The estimated size of each PaaS database is 1 TB.
Contoso plans to implement the following changes:
Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.
Start onboarding customers to the new PaaS solution within six months.
Business Goals
Contoso identifies the following business requirements:
Use built-in Azure features whenever possible.
Minimize development effort whenever possible.
Minimize the compute costs of the PaaS solutions.
Provide all the customers with their own copy of the database by using the PaaS solution.
Provide the customers with different table and row access based on the customer's service agreement. In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover. Ensure that users of the PaaS solution can create their own database objects but 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.

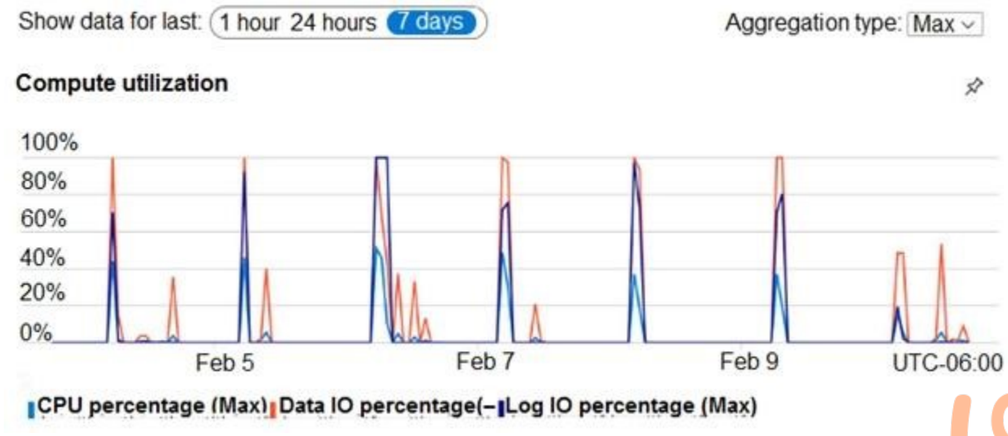

Role Assignments
For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

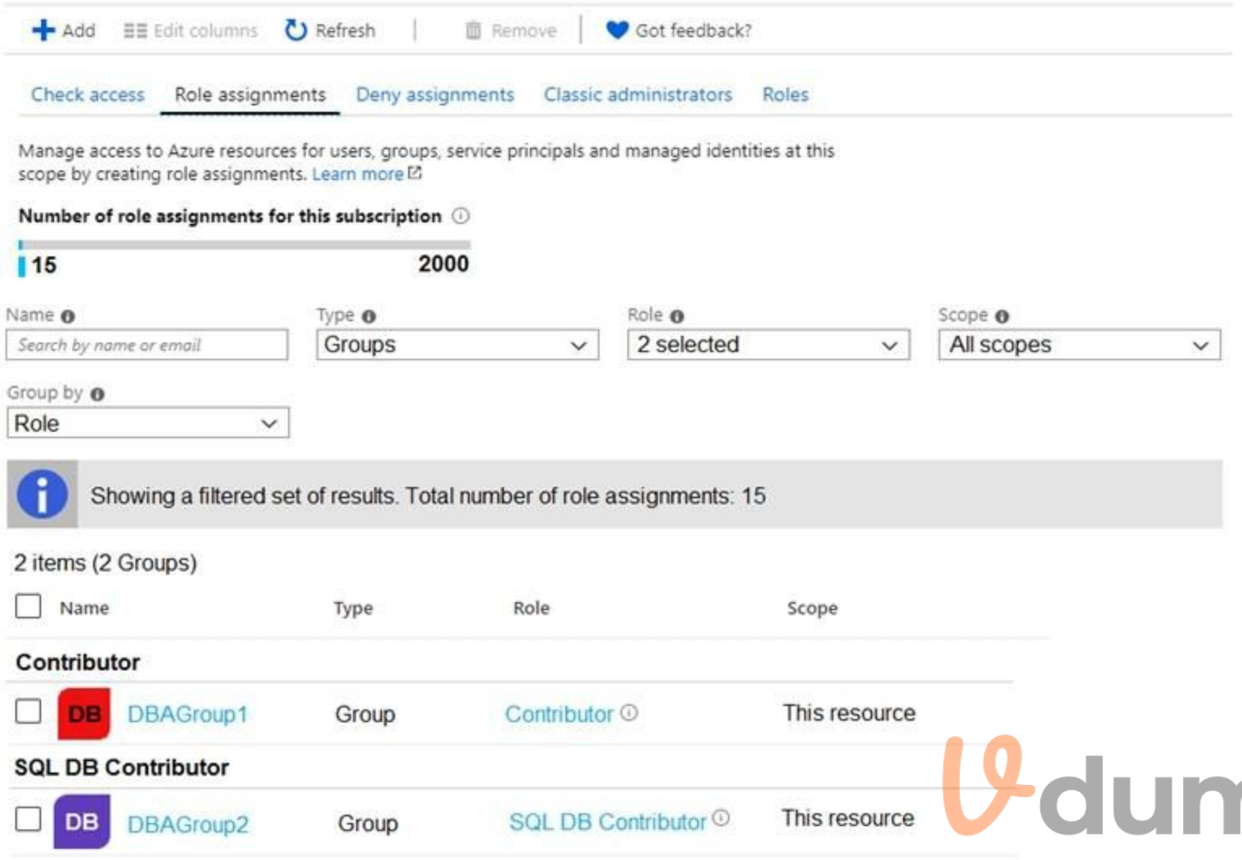

QUESTION 1
You need to recommend a solution to ensure that the customers can create the database objects. The solution must meet the business goals. What should you include in the recommendation?

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.

Correct Answer: B
Section:

QUESTION 2
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)

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-ver15

QUESTION 3
HOTSPOT
You are evaluating the role assignments.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Hot Area:

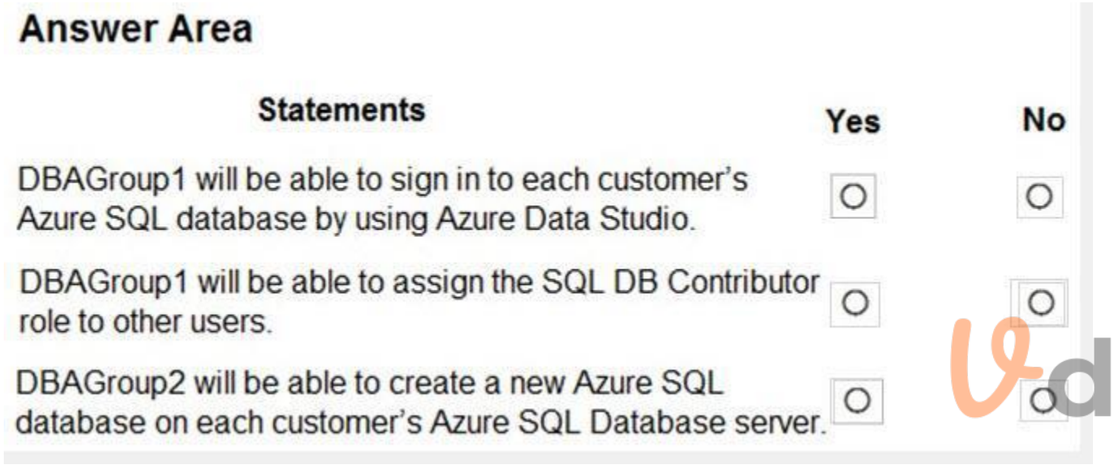

Answer Area:

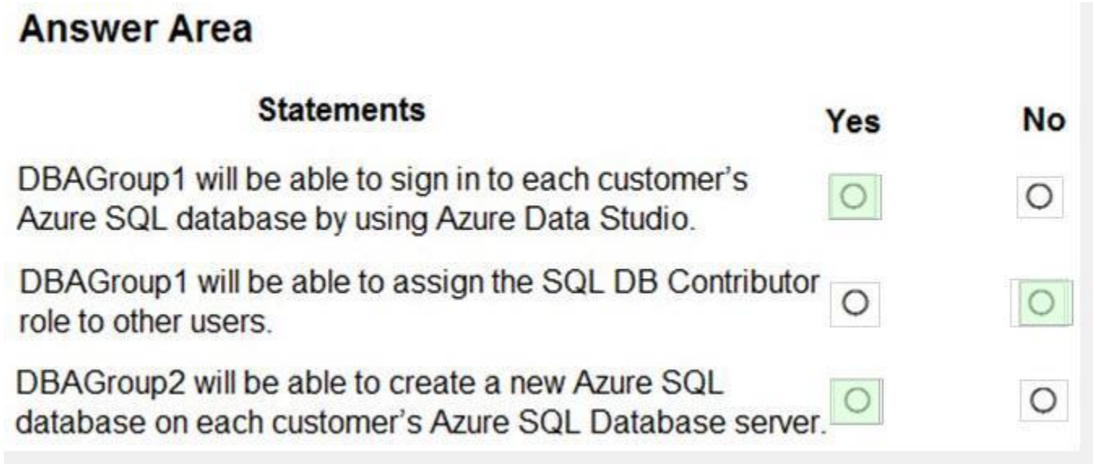

Section:
Explanation:
Box 1: Yes
DBAGroup1 is member of the Contributor role.
The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.
Box 2: No

Box 3: Yes
DBAGroup2 is member of the SQL DB Contributor role.
The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.
Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

03 - Implement a Secure Environment

QUESTION 1
You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information. You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information. What should you include in the recommendation?

A. row-level security
B. data masking
C. Always Encrypted
D. column-level security

Correct Answer: B
Section:
Explanation:
Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users. The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card. Example:
XXXX-XXXX-XXXX-1234

QUESTION 2
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

Correct Answer: A
Section:
Explanation:
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

QUESTION 3
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

Correct Answer: D, E
Section:
Explanation:
Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes. Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.
Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline

QUESTION 4
You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. DF1 contains a linked service. You have an Azure Key vault named vault1 that contains an encryption kay named key1.
You need to encrypt df1 by using key1.
What should you do first?

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

Correct Answer: B
Section:
Explanation:
A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.
Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.
Incorrect Answers:
A, D: Should enable Soft Delete and Do Not Purge on Azure Key Vault.
Using customer-managed keys with Data Factory requires two properties to be set on the Key Vault, Soft Delete and Do Not Purge. These properties can be enabled using either PowerShell or Azure CLI on a new or existing key vault.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key

QUESTION 5
You have a new Azure SQL database. The database contains a column that stores confidential information. You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed. Which three actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Turn on auditing and write audit logs to an Azure Storage account.
B. Add extended properties to the column.
C. Turn on Advanced Data Security for the Azure SQL server.
D. Apply sensitivity labels named Highly Confidential to the column.
E. Turn on Azure Advanced Threat Protection (ATP).

Correct Answer: A, C, D

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

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

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

QUESTION 6
You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.
You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.
You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:
Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.
What should you create on VNet1?

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

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

QUESTION 7
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.

Correct Answer: B, E
Section:
Explanation:
By default, the cluster of nodes for the premium availability model is created in the same datacenter.With the introduction of Azure Availability Zones, SQL Database can place different replicas of theBusiness Critical database to different availability zones in the same region. To eliminate a singlepoint of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zoneredundant configuration in the Premium or Business Critical service tiers does not create additionaldatabase redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration,you can make your Premium or Business Critical databases resilient to a much larger set of failures,including catastrophic datacenter outages, without any changes to the application logic. You can alsoconvert any existing Premium or Business Critical databases or pools to the zone redundantconfiguration.

QUESTION 8
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

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

QUESTION 9
HOTSPOT
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.

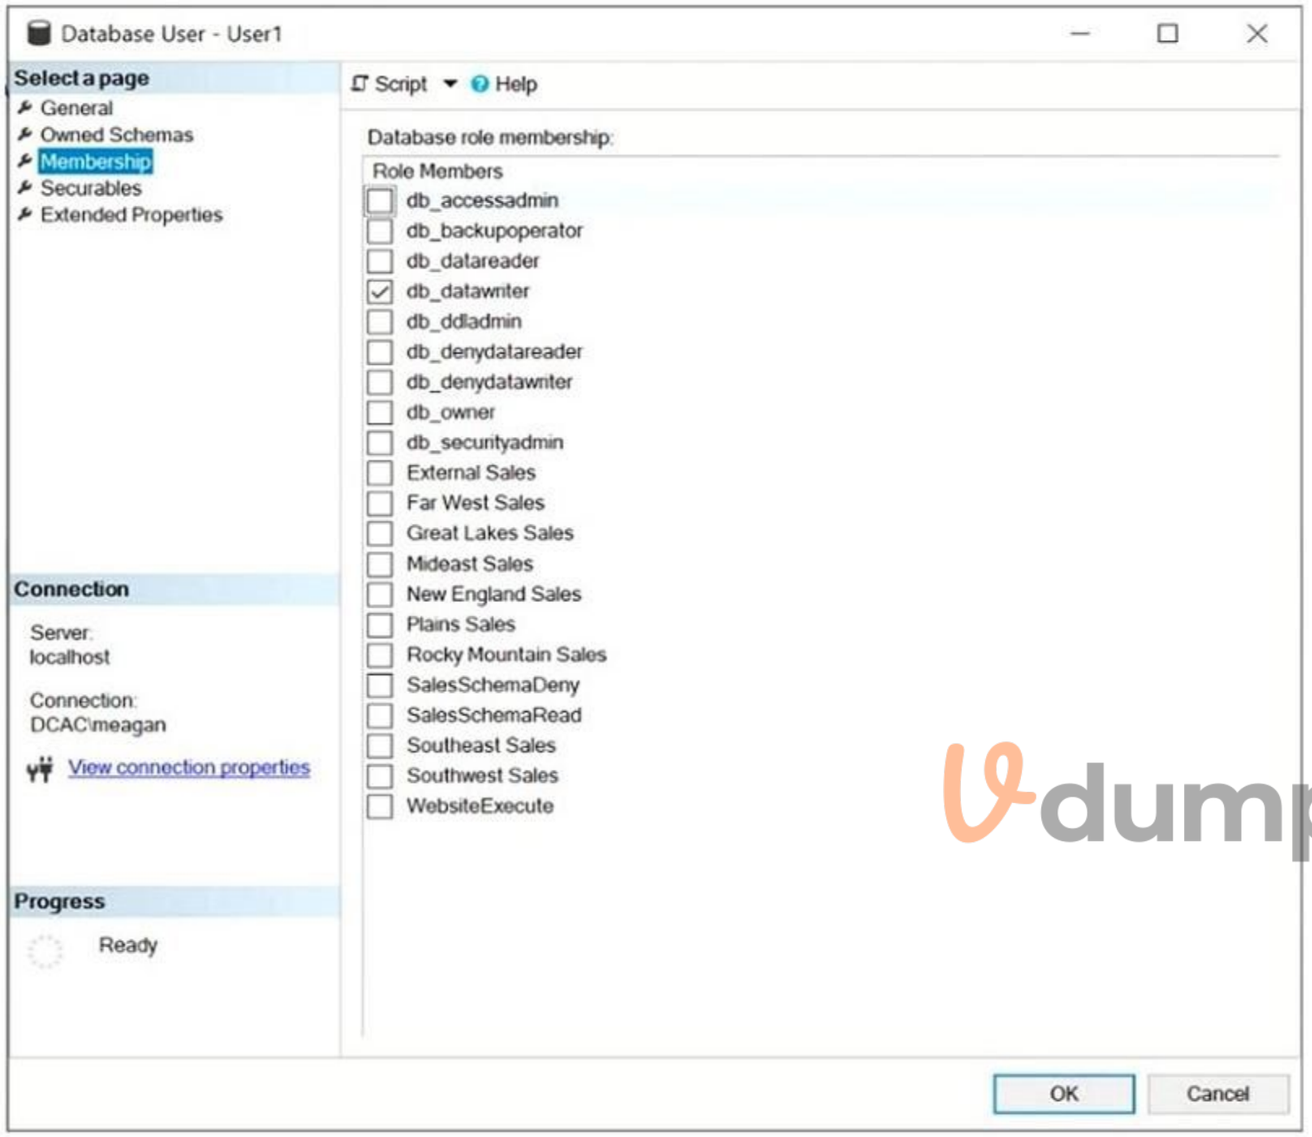

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

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 Area:

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

Section:
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

QUESTION 10
DRAG DROP
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.

Select and Place:

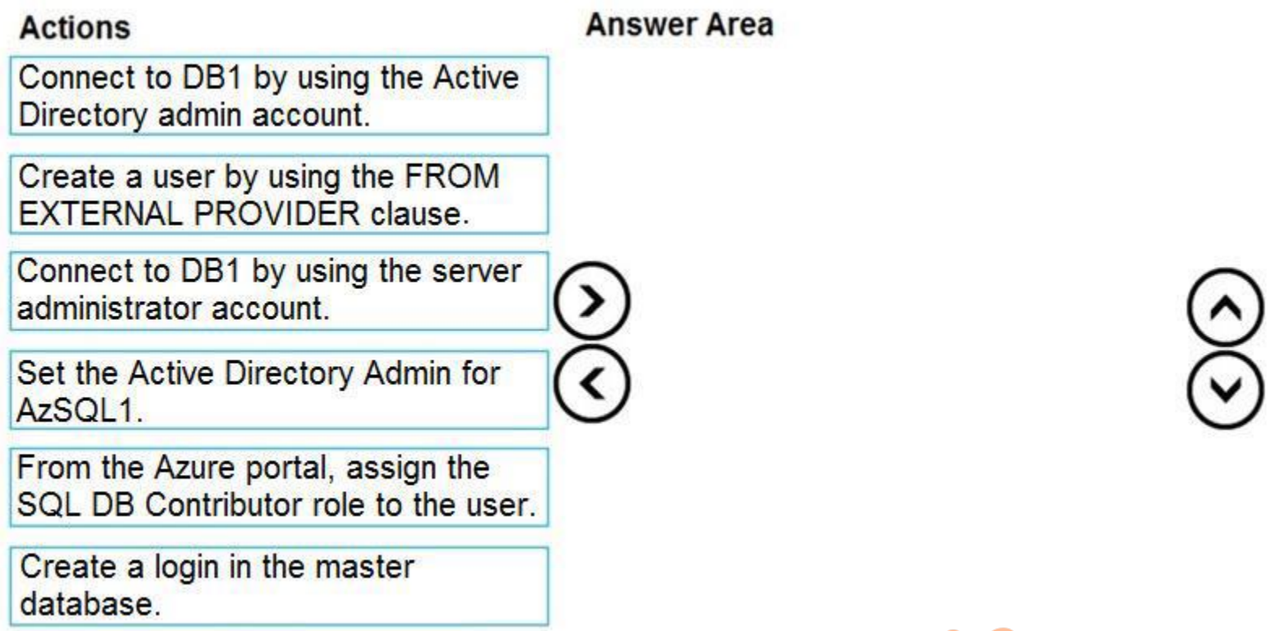

Correct Answer:

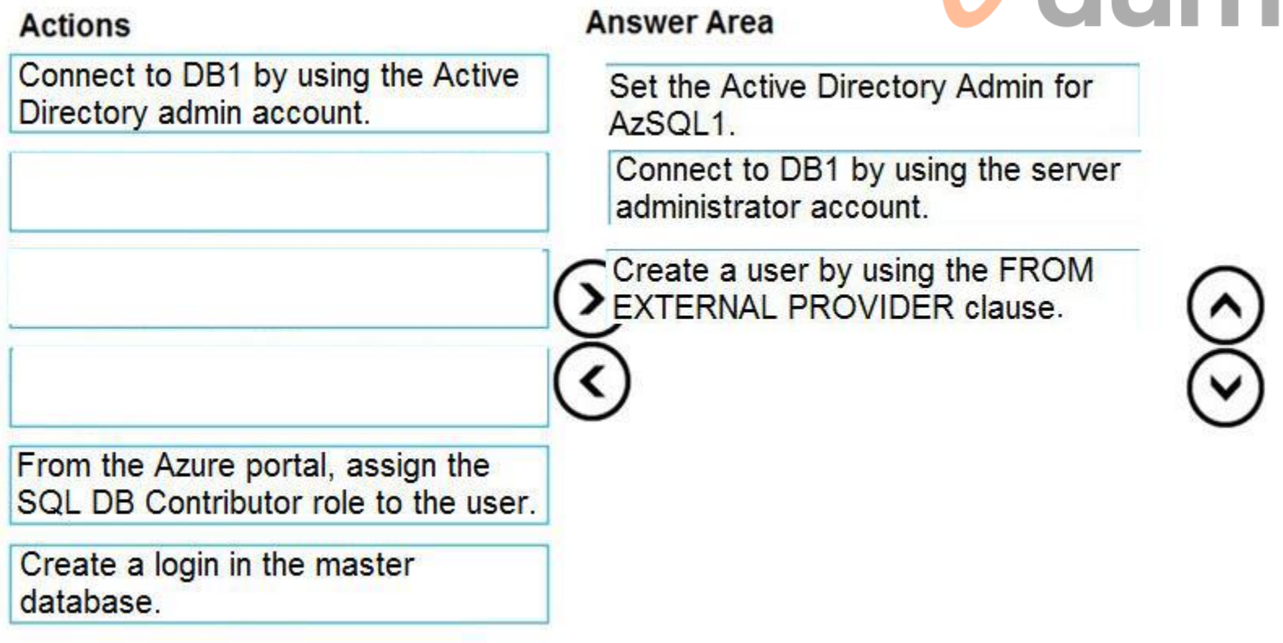

Section:
Explanation:
Step 1: Set up the Active Directory Admin for AzSQL1.
Step 2: Connect to DB1 by using the server administrator.

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

QUESTION 11
HOTSPOT
You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

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

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:
- The first six numerals of each customer's phone number must be masked.
- The first six numerals of each customer's phone number must be masked.
- The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.
How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Hot Area:

Answer Area

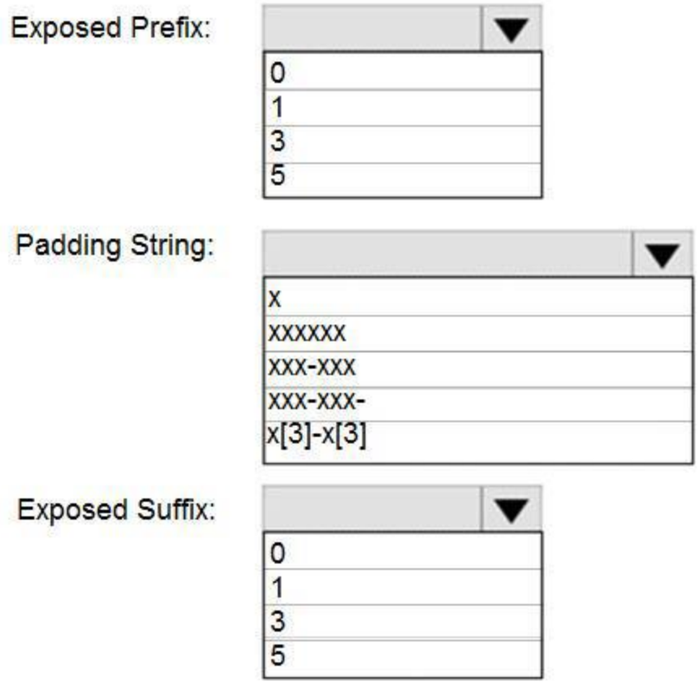

Answer Area:

Section:

Explanation:
Box 1: 0
Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix
Box 2: xxx-xxx
Box 3: 5
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking

QUESTION 12
DRAG DROP
You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.
You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

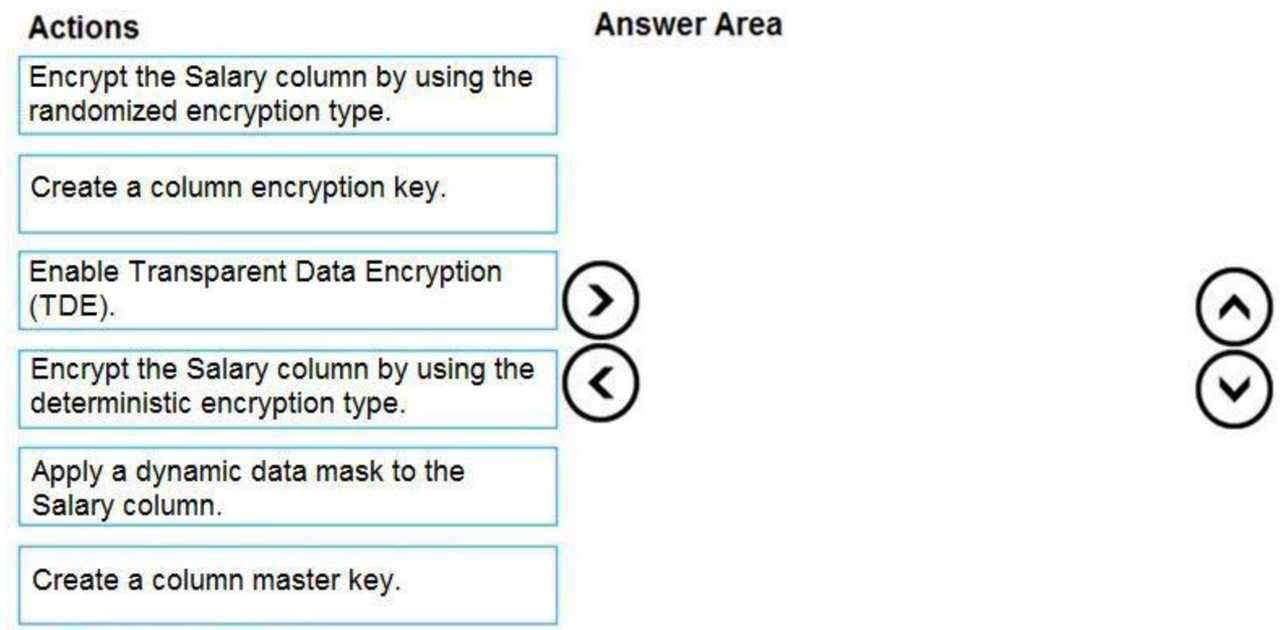

Correct Answer:

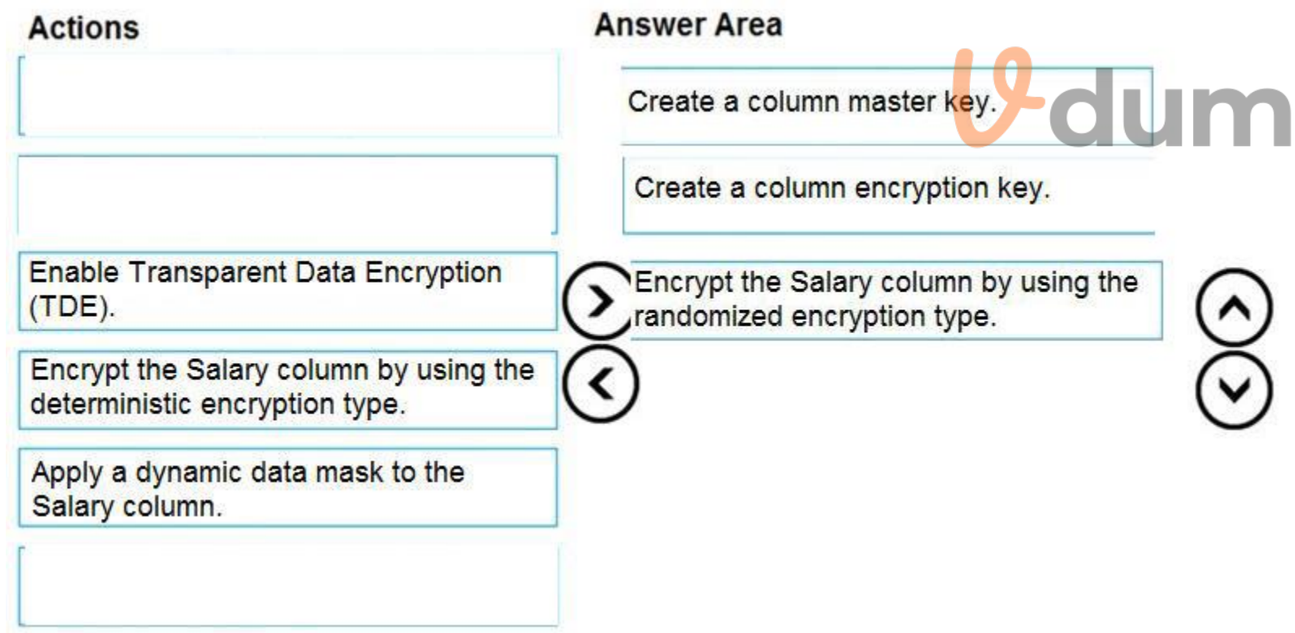

Section:
Explanation:
Step 1: Create a column master key
Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.
Step 2: Create a column encryption key.
Step 3: Encrypt the Salary column by using the randomized encryption type.
Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.
Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Incorrect Answers:
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. However, it may also allow unauthorized users to guess information about encrypted values by examining patterns in the encrypted column, especially if there's a small set of possible encrypted values, such as True/False, or North/South/East/West region.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine

QUESTION 13
HOTSPOT
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.

Hot Area:

Answer Area

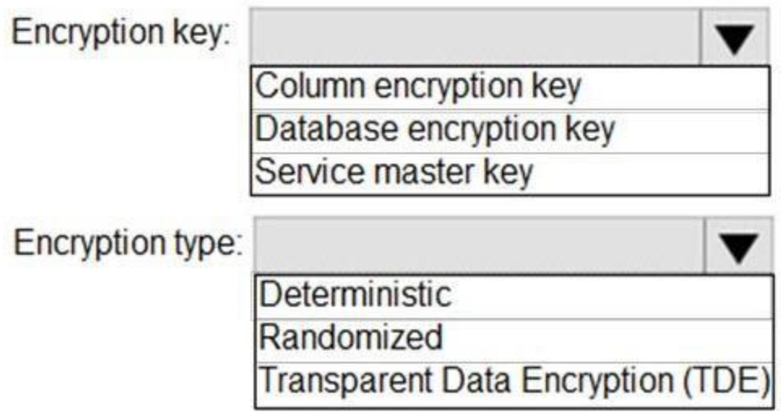

Answer Area:

Answer Area

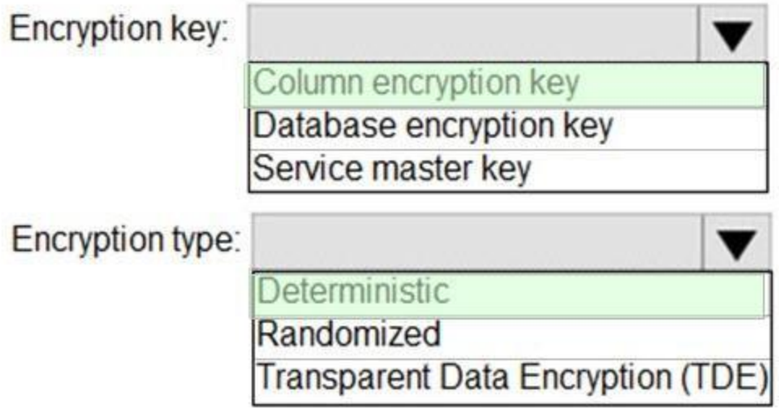

Section:
Explanation:
Box 1: Column encryption Key
Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.
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

QUESTION 14
DRAG DROP
You have an Azure SQL Database instance named DatabaseA on a server named Server1.
You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.
You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

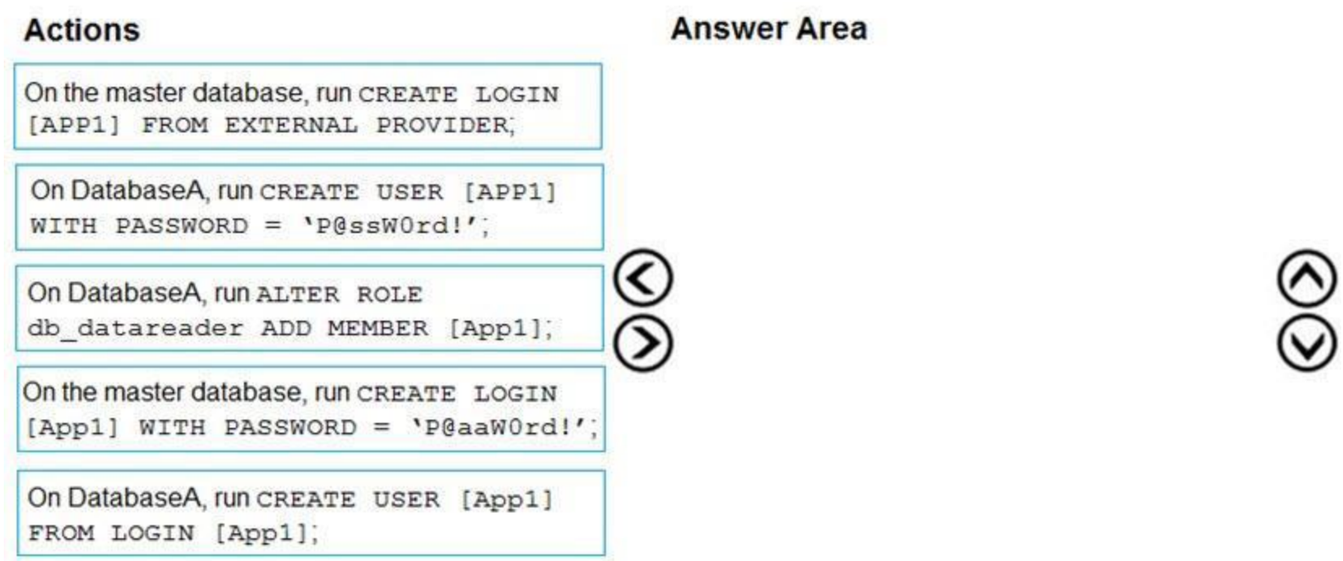

Correct Answer:

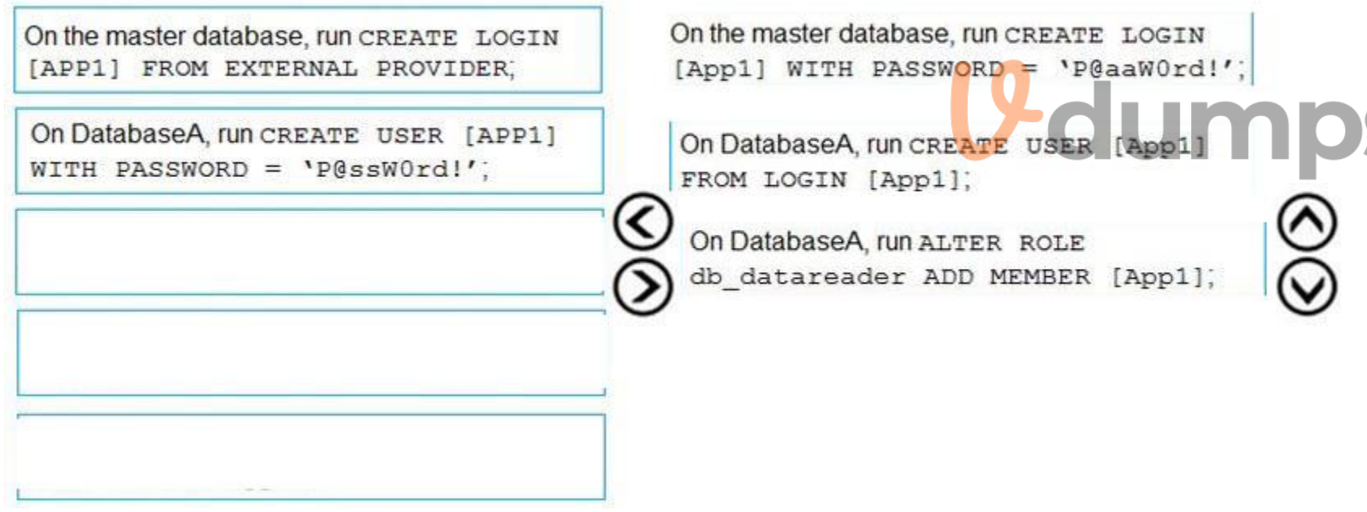

Section:
Explanation:
Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:
CREATE LOGIN readonlylogin WITH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

db_datareader role.
EXEC sp_addrolemember 'db_datareader', 'readonlyuser';
Reference:
https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/

QUESTION 15
You are developing an application that uses Azure Data Lake Storage Gen 2.
You need to recommend a solution to grant permissions to a specific application for a limited time period. What should you include in the recommendation?

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

Correct Answer: C
Section:
Explanation:
A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:
What resources the client may access.
What permissions they have to those resources.
How long the SAS is valid.
Note: Data Lake Storage Gen2 supports the following authorization mechanisms:
Shared Key authorization
Shared access signature (SAS) authorization
Role-based access control (Azure RBAC)
Access control lists (ACL) Data Lake Storage Gen2 supports the following authorization mechanisms:
Shared Key authorization
Shared access signature (SAS) authorization
Role-based access control (Azure RBAC)
Access control lists (ACL)
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

QUESTION 16
HOTSPOT
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.

Hot Area:

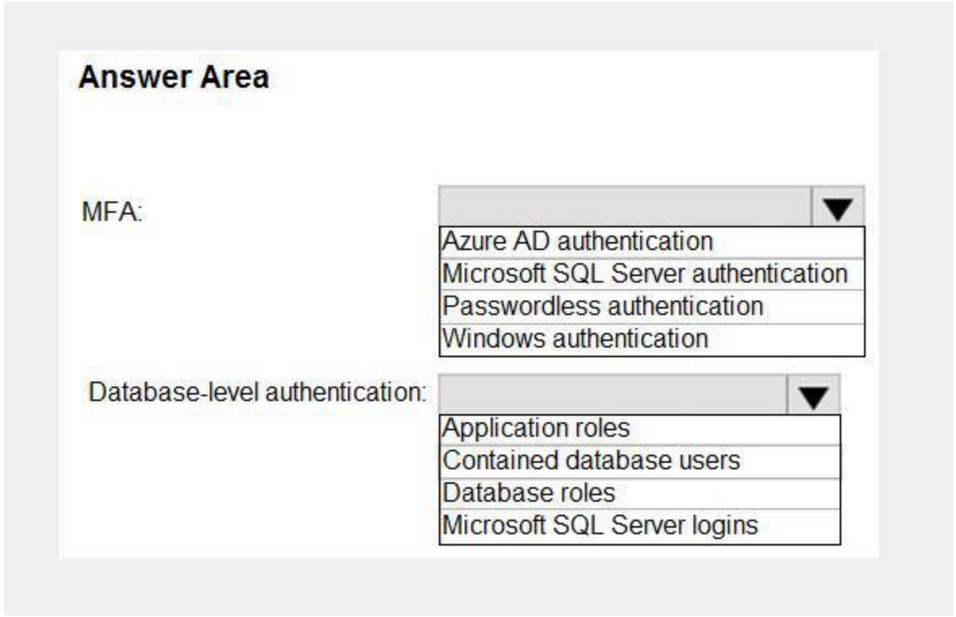

Answer Area:

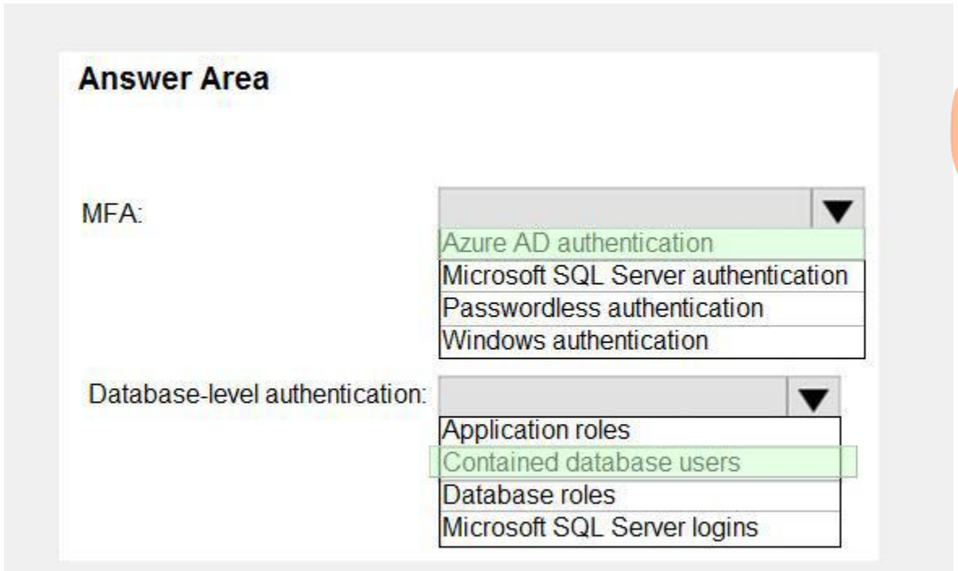

Section:
Explanation:
Box 1: Azure AD authentication
Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.
Box 2: Contained database users
Azure Active Directory Uses contained database users to authenticate identities at the database 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

QUESTION 17
You have an Azure subscription that contains a server named Server1. Server1 hosts two Azure SQL databases named DB1 and DB2.
You plan to deploy a Windows app named App1 that will authenticate to DB2 by using SQL authentication.
You need to ensure that App1 can access DB2. The solution must meet the following requirements:
App1 must be able to view only DB2.
Administrative effort must be minimized.
What should you create?

A. a contained database user for App1 on DB2
B. a login for App1 on Server1
C. a contained database user from an external provider for App1 on DB2
D. a contained database user from a Windows login for App1 on DB2

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/security/contained-database-users-making-your-database-portable?view=sql-server-ver15

QUESTION 18
You create five Azure SQL Database instances on the same logical server.
In each database, you create a user for an Azure Active Directory (Azure AD) user named User1.
User1 attempts to connect to the logical server by using Azure Data Studio and receives a login error.
You need to ensure that when User1 connects to the logical server by using Azure Data Studio, User1 can see all the databases.
What should you do?

A. Create User1 in the master database.
B. Assign User1 the db_datareader role for the master database.
C. Assign User1 the db_datareader role for the databases that User1 creates.
D. Grant SELECT on sys.databases to public in the master database.

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/logins-create-manage

01 - Optimize Query Performance
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot

return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
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.

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

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

Correct Answer: A, E
Section:
Explanation:
SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.
A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.
E: Example:
The following query will show blocking information.
SELECT
t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
Note: Depending on the system you're working with you can access these wait statistics from one of three locations:
sys.dm_os_wait_stats: for SQL Server
sys.dm_db_wait_stats: for Azure SQL Database
sys.dm_pdw_nodes_os_wait_stats: for Azure SQL Data Warehouse
Incorrect Answers:
F: sys.dm_tran_locks returns information about currently active lock manager resources in SQL Server 2019 (15.x). Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.
Instead use sys.dm_pdw_nodes_tran_locks.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql

02 - Optimize Query Performance

QUESTION 1
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

Correct Answer: C
Section:
Explanation:
Use the Query Store Page in SQL Server Management Studio.
Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve. Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store

QUESTION 2
You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently. You enable columnstore archival compression.
What are two possible results of the configuration? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Queries that use the index will consume more disk I/O.
B. Queries that use the index will retrieve fewer data pages.
C. The index will consume more disk space.
D. The index will consume more memory.
E. Queries that use the index will consume more CPU resources.

Correct Answer: B, E
Section:
Explanation:
For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.
Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

QUESTION 3
HOTSPOT
You have SQL Server on an Azure virtual machine.
You review the query plan shown in the following exhibit.

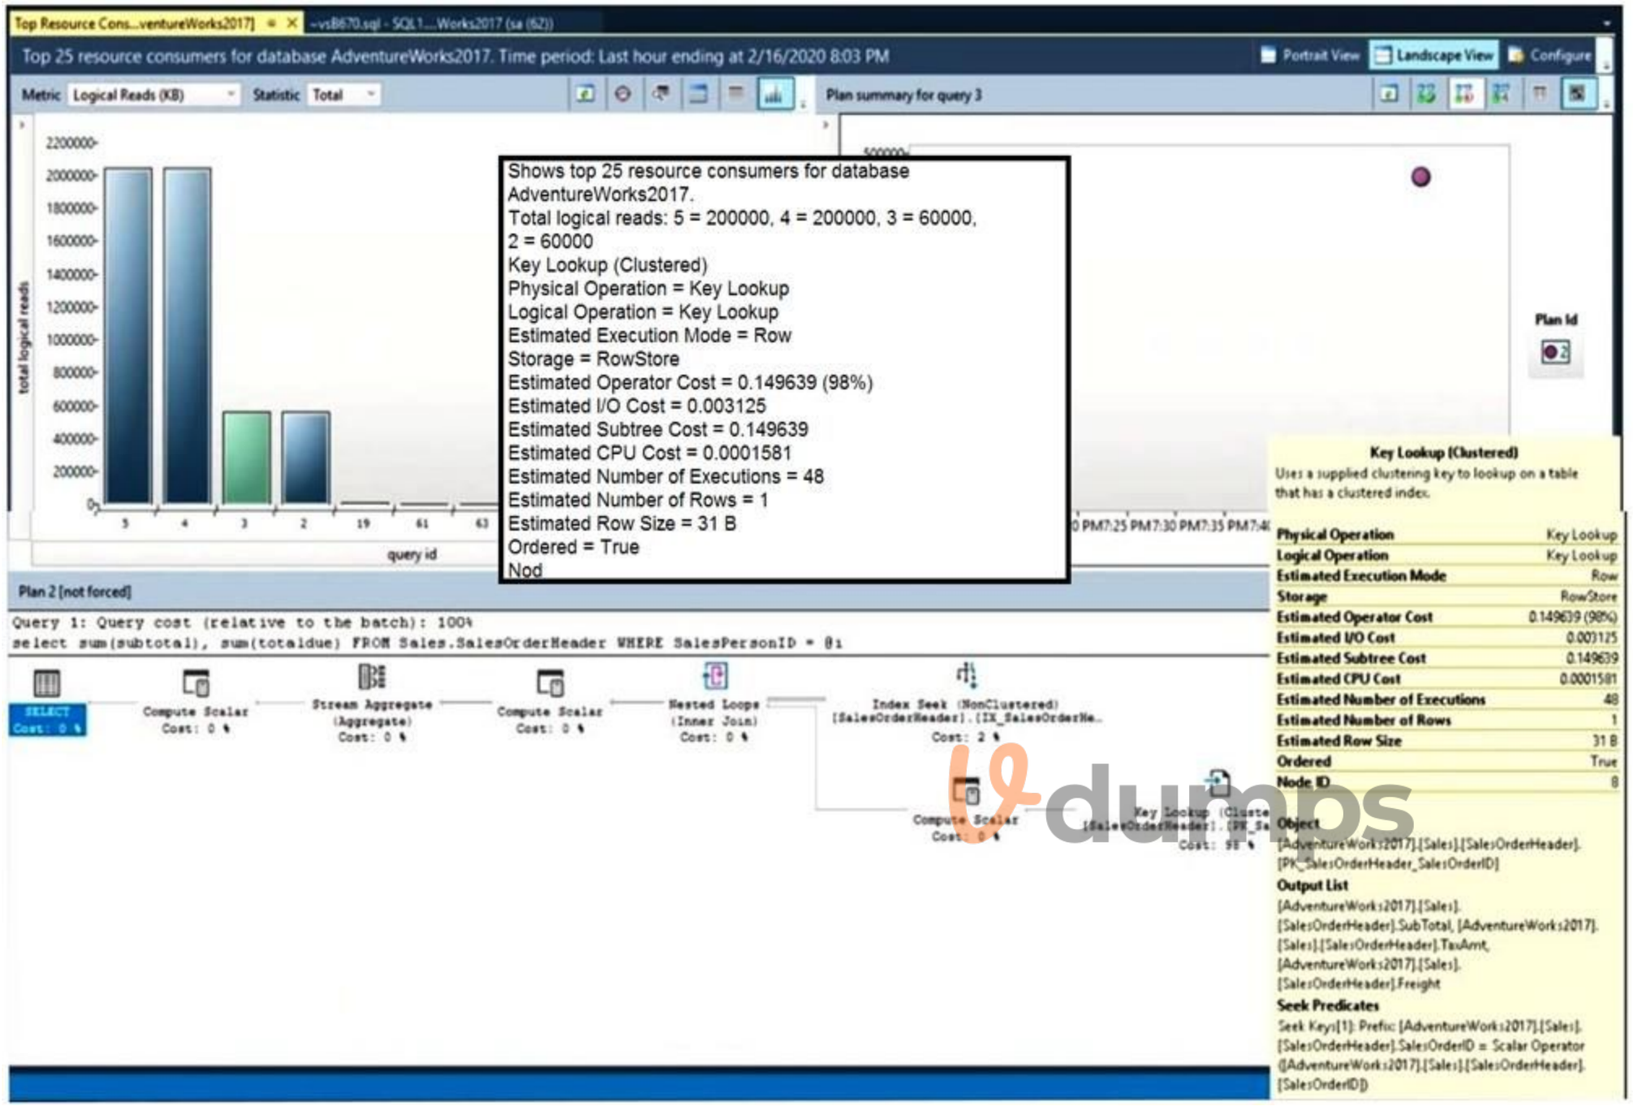

For each of the following statements, select yes if the statement is true. Otherwise, select no.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	○	○
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	○	○
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	○	○

Answer Area:

Answer Area

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	○	○
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	○	○
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	○	○

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

QUESTION 4
HOTSPOT
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.

Hot Area:

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 Area:

Answer Area

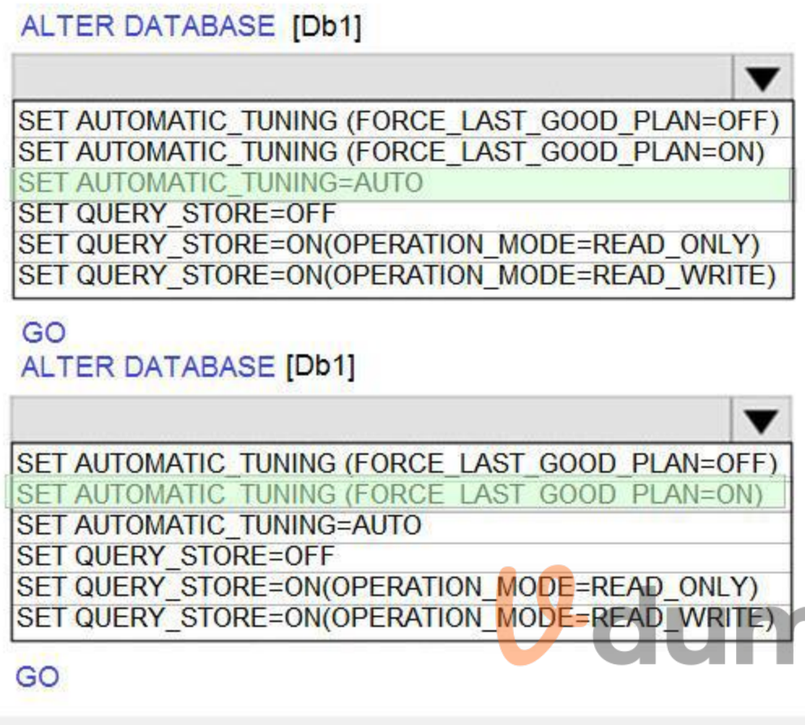

Section:
Explanation:
Box 1: SET AUTOMATIC_TUNING = AUTO
To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:
ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO
Setting automatic tuning to AUTO will apply Azure Defaults.
Box 2: SET AUTOMATIC_TUNING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

QUESTION 5
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.

Correct Answer: A, E
Section:

QUESTION 6
You have an Azure SQL database.
You identify a long running query.
You need to identify which operation in the query is causing the performance issue.
What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

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

Correct Answer: A
Section:
Explanation:
https://www.mssqltips.com/sqlservertip/3685/live-query-statistics-in-sql-server-2016/

QUESTION 7
You have a version-8.0 Azure Database for MySQL database.
You need to identify which database queries consume the most resources.
Which tool should you use?

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

Correct Answer: A
Section:
Explanation:
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

QUESTION 8
You have SQL Server on an Azure virtual machine that contains a database named DB1.
You have an application that queries DB1 to generate a sales report.
You need to see the parameter values from the last time the query was executed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Enable Last_Query_Plan_Stats in the master database
B. Enable Lightweight_Query_Profiling in DB1
C. Enable Last_Query_Plan_Stats in DB1
D. Enable Lightweight_Query_Profiling in the master database
E. Enable PARAMETER_SNIFFING in DB1

Correct Answer: B, E
Section:
Explanation:

QUESTION 9
You deploy a database to an Azure SQL Database managed instance.
You need to prevent read queries from blocking queries that are trying to write to the database.
Which database option should set?

A. PARAMETERIZATION to FORCED
B. PARAMETERIZATION to SIMPLE
C. Delayed Durability to Forced
D. READ_COMMITTED_SNAPSHOT to ON

Correct Answer: D
Section:
Explanation:
In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:
The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.
The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.
Incorrect Answers:
A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may choose to parameterize the queries. This means that any literal values that are contained in a query are substituted with parameters. This process is referred to as simple parameterization. When SIMPLE parameterization is in effect, you cannot control which queries are parameterized and which queries are not.
B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION database option to FORCED. This process is referred to as forced parameterization.
C: Delayed transaction durability is accomplished using asynchronous log writes to disk. Transaction log records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place. Delayed transaction durability reduces both latency and contention within the system.
Some of the cases in which you could benefit from using delayed transaction durability are:
You can tolerate some data loss.
You are experiencing a bottleneck on transaction log writes.
Your workloads have a high contention rate.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql

QUESTION 10
You have an Azure SQL database.
You discover that the plan cache is full of compiled plans that were used only once.
You run the select * from sys.database_scoped_configurations Transact-SQL command and receive the results shown in the following table.

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

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

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

Correct Answer: C
Section:
Explanation:
OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }
Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.
Incorrect Answers:
A: LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }
Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and earlier version independent of the compatibility level of the database. The default is OFF, which sets the query optimizer cardinality estimation model based on the compatibility level of the database.
B: QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY }
Enables or disables query optimization hotfixes regardless of the compatibility level of the database. The default is OFF, which disables query optimization hotfixes that were released after the highest available compatibility level was introduced for a specific version (post-RTM).
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql

QUESTION 11

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:

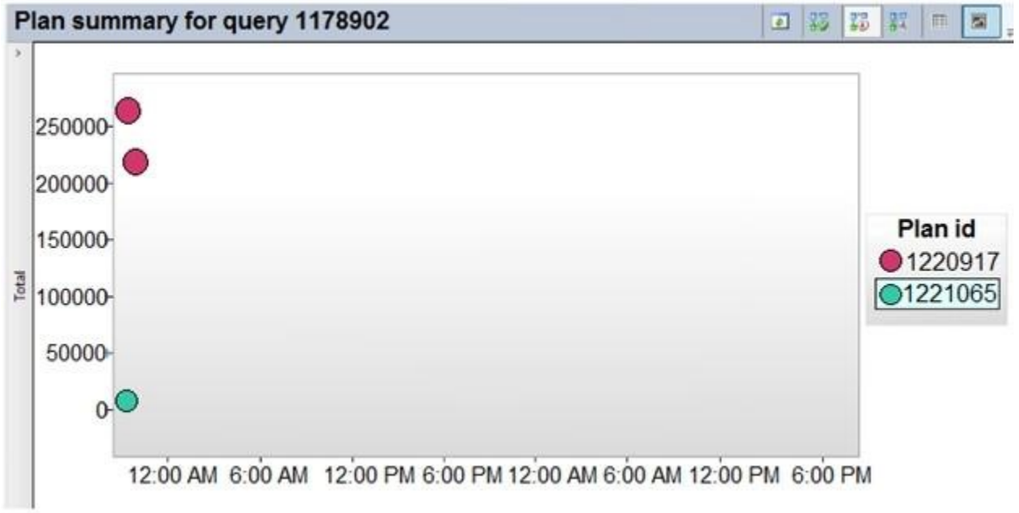

What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

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.

Correct Answer: C
Section:
Explanation:
Reference: https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios

QUESTION 12
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.

Correct Answer: C
Section:
Explanation:
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.

QUESTION 13
HOTSPOT
You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.

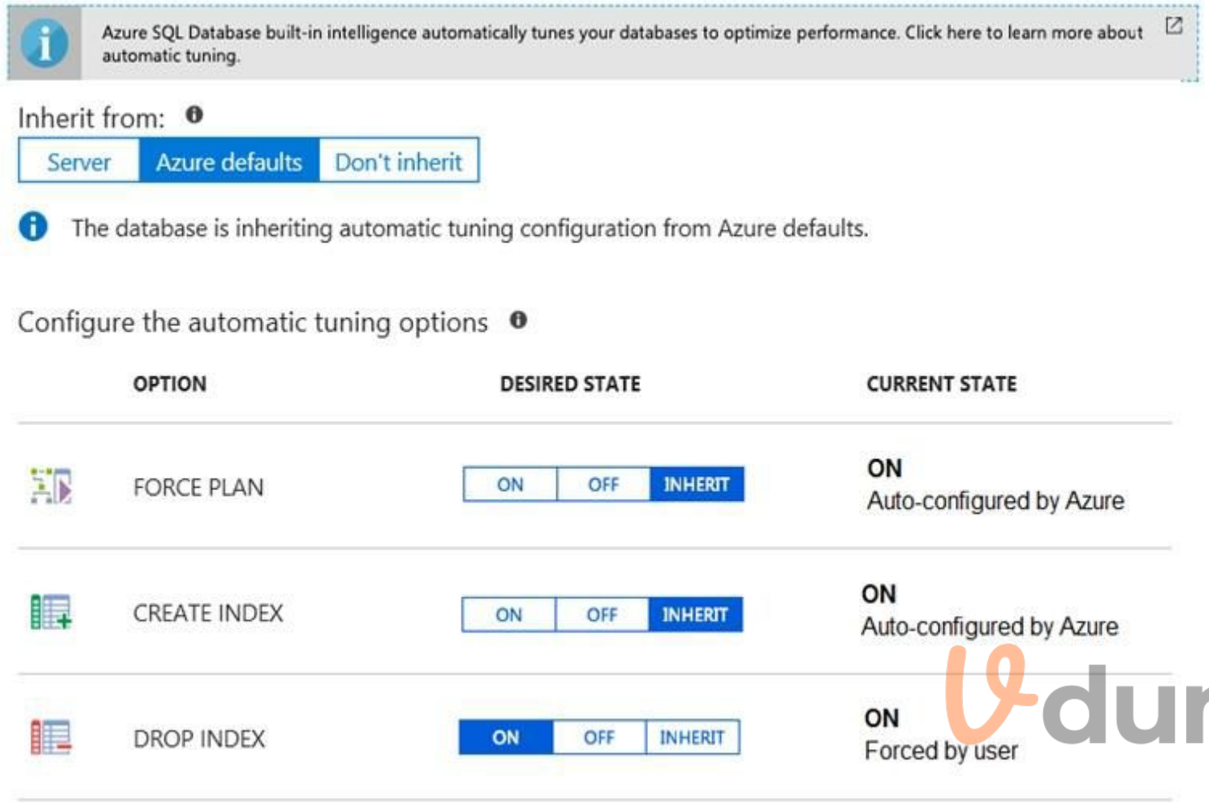

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Hot Area:

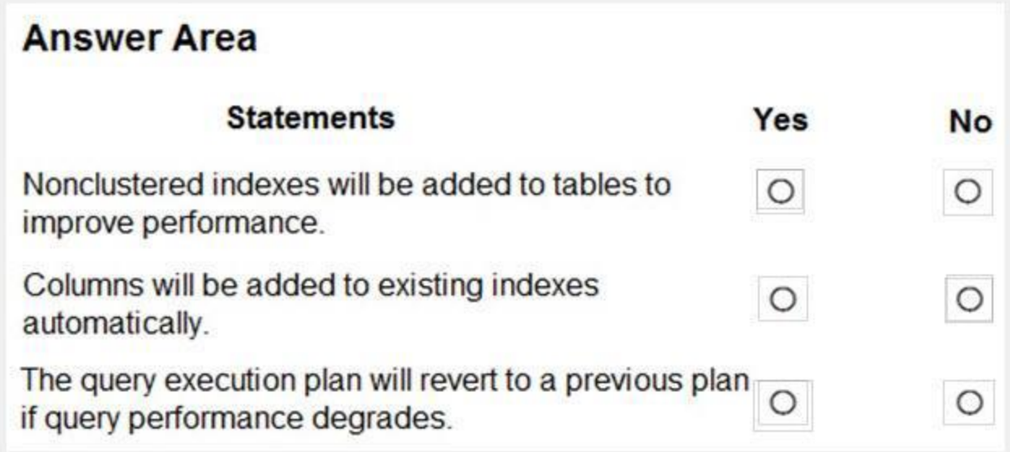

Answer Area:

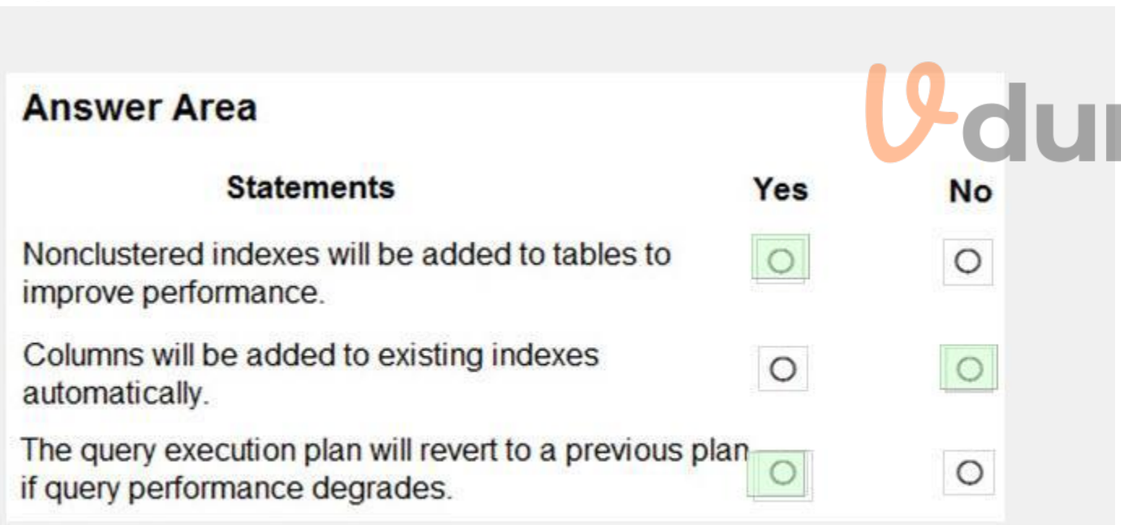

Section:
Explanation:
Box 1: Yes
We see: Tuning option: Create index ON
CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.
Box 2: No
Box 3: Yes
FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

QUESTION 14
You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.
You need to create a surrogate key for the table. The solution must provide the fastest query performance. What should you use for the surrogate key?

A. an IDENTITY column
B. a GUID column
C. a sequence object

Correct Answer: A
Section:
Explanation:
Dedicated SQL pool supports many, but not all, of the table features offered by other databases. Surrogate keys are not supported. Implement it with an Identity column.
Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-overview

QUESTION 15
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.

Correct Answer: B, D
Section:
Explanation:
Reference: https://community.idera.com/database-tools/blog/b/community_blog/posts/why-use-a-date-dimension-table-in-a-data-warehouse

QUESTION 16
You have an Azure Data Factory pipeline that is triggered hourly.
The pipeline has had 100% success for the past seven days.
The pipeline execution fails, and two retries that occur 15 minutes apart also fail. The third failure returns the following error.

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

What is a possible cause of the error?

A. From 06:00 to 07:00 on January 10, 2021, there was no data in wwi/BIKES/CARBON.
B. The parameter used to generate year=2021/month=01/day=10/hour=06 was incorrect.

C. From 06:00 to 07:00 on January 10, 2021, the file format of data in wwi/BIKES/CARBON was incorrect.
D. The pipeline was triggered too early.

Correct Answer: B
Section:

QUESTION 17
HOTSPOT
You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.
You plan to have a fact table for website visits. The table will be approximately 5 GB.
You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

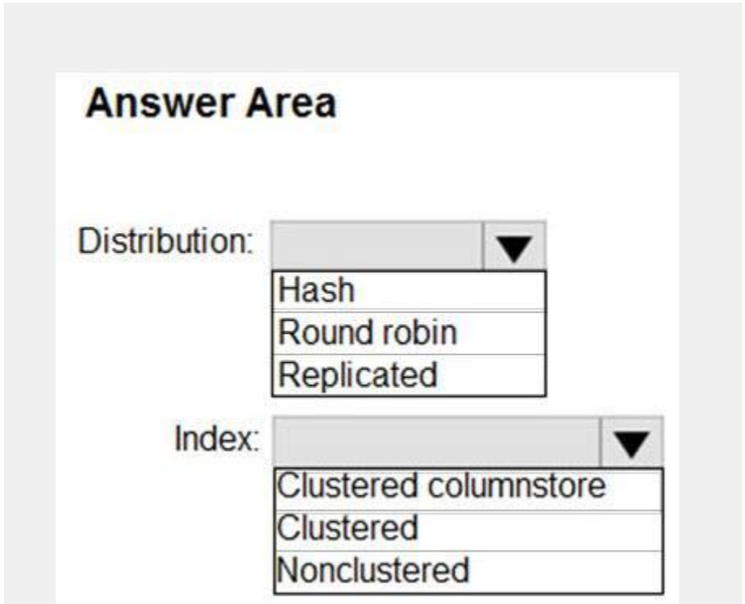

Answer Area:

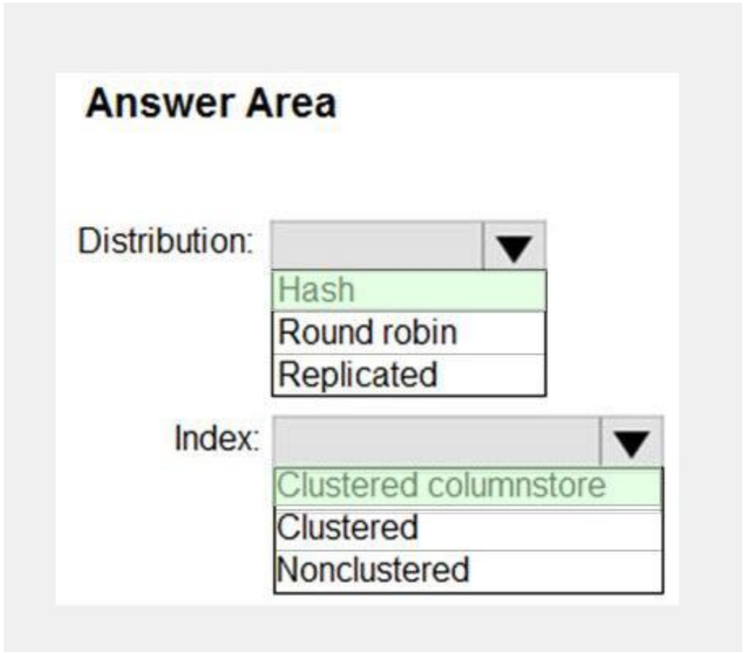

Section:
Explanation:
Box 1: Hash
Consider using a hash-distributed table when:
The table size on disk is more than 2 GB.
The table has frequent insert, update, and delete operations.
Box 2: Clustered columnstore
Clustered columnstore tables offer both the highest level of data compression and the best overall query performance.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-index

QUESTION 18
HOTSPOT
You have an Azure SQL database.
You are reviewing a slow performing query as shown in the following exhibit.

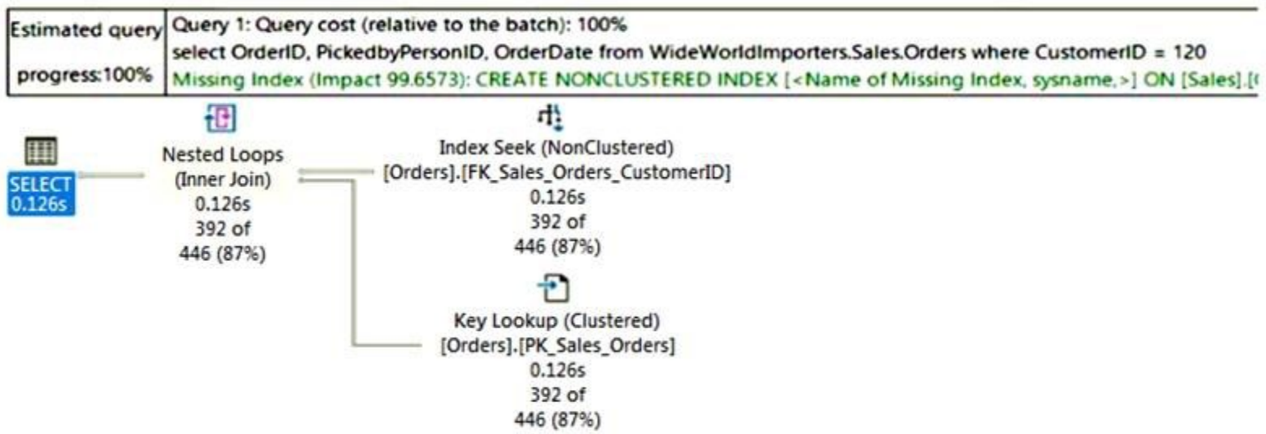

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

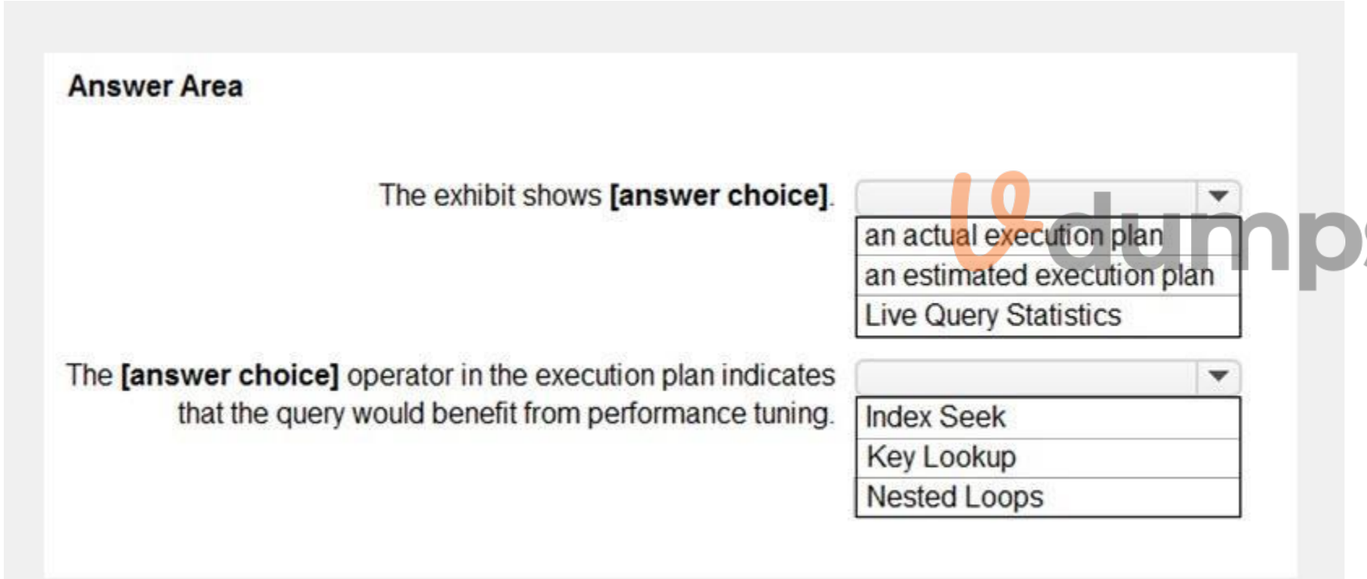

Answer Area:

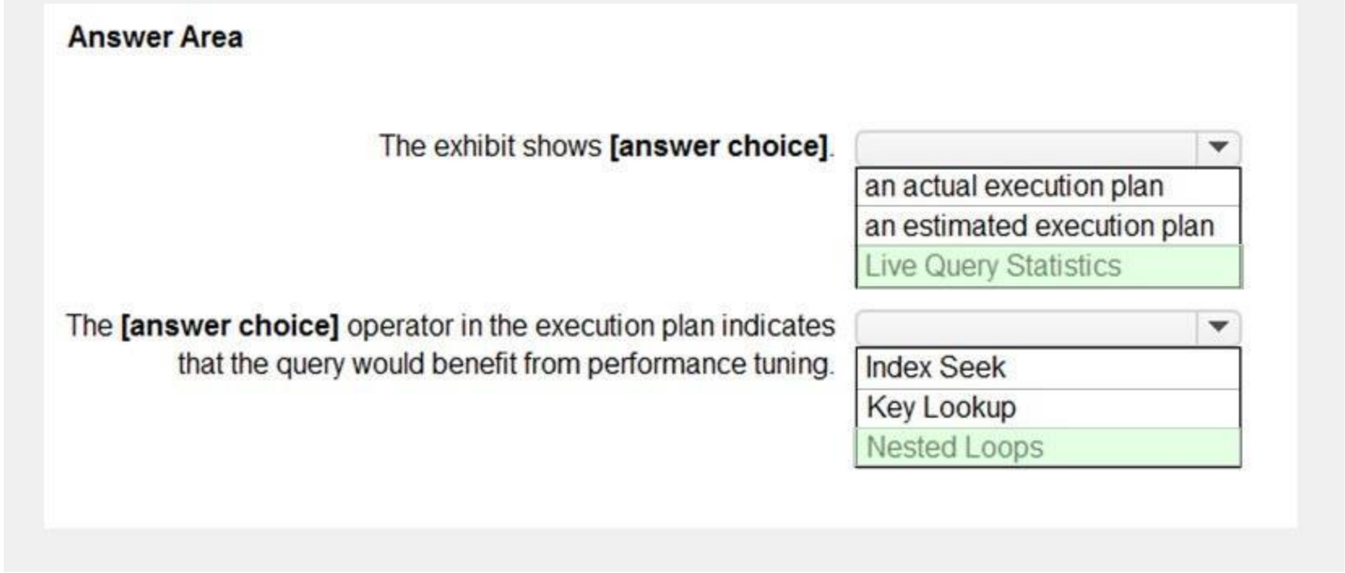

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-ver15

QUESTION 19
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.

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15

QUESTION 20
HOTSPOT
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.

Hot Area:

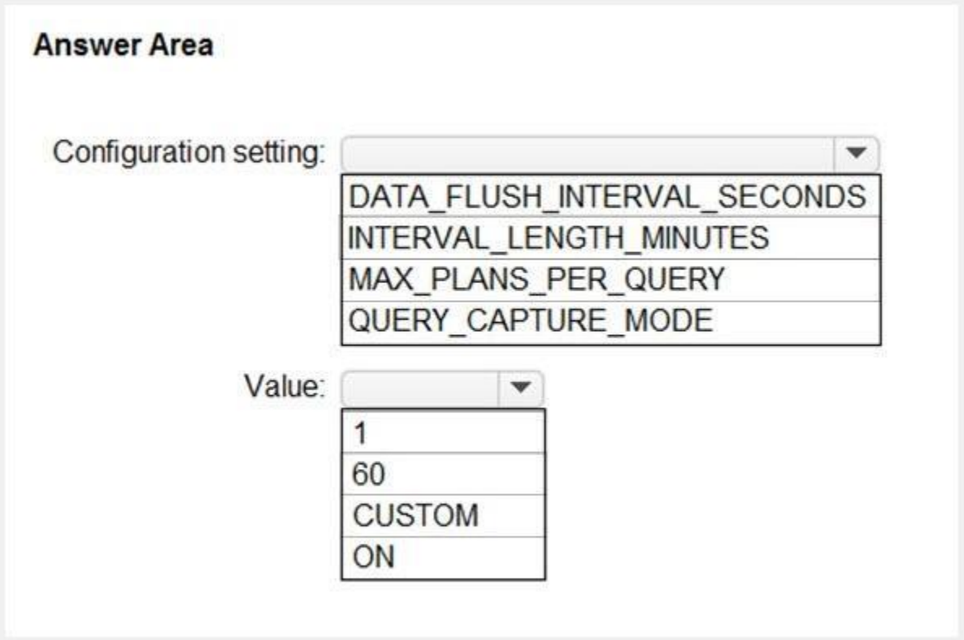

Answer Area:

Answer Area

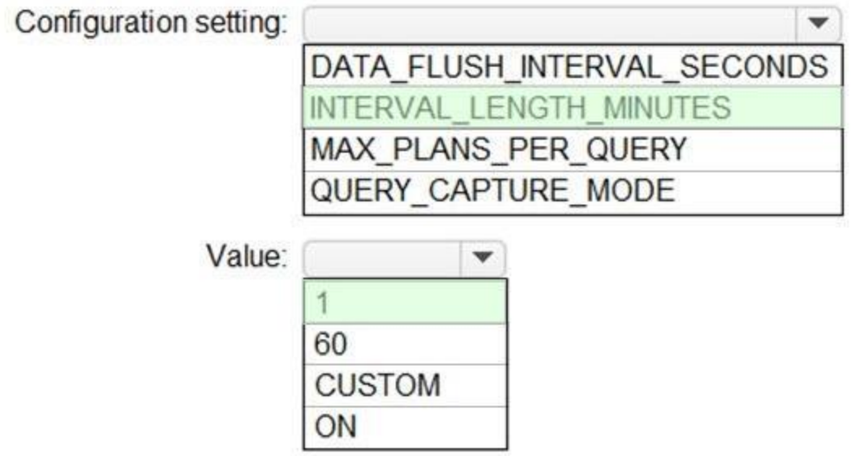

Section:
Explanation:

QUESTION 21
You have an Azure SQL Database managed instance.
The instance starts experiencing performance issues.
You need to identify which query is causing the issue and retrieve the execution plan for the query. The solution must minimize administrative effort.
What should you use?

A. SQL Profiler
B. Extended Events
C. Query Store
D. dynamic management views

Correct Answer: D
Section:

QUESTION 22
You have an Azure SQL database named DB1.
You need to display the estimated execution plan of a query by using the query editor in the Azure portal.
What should you do first?

A. Run the SET SHOWPLAN_ALL Transact-SQL statement.

B. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
C. Run the SET FORCEPLAN Transact-SQL statement.
D. Enable Query Store for DB1.

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

QUESTION 23
HOTSPOT
You have an Azure SQL database.
You have a query and the associated execution plan as shown in the following exhibit.

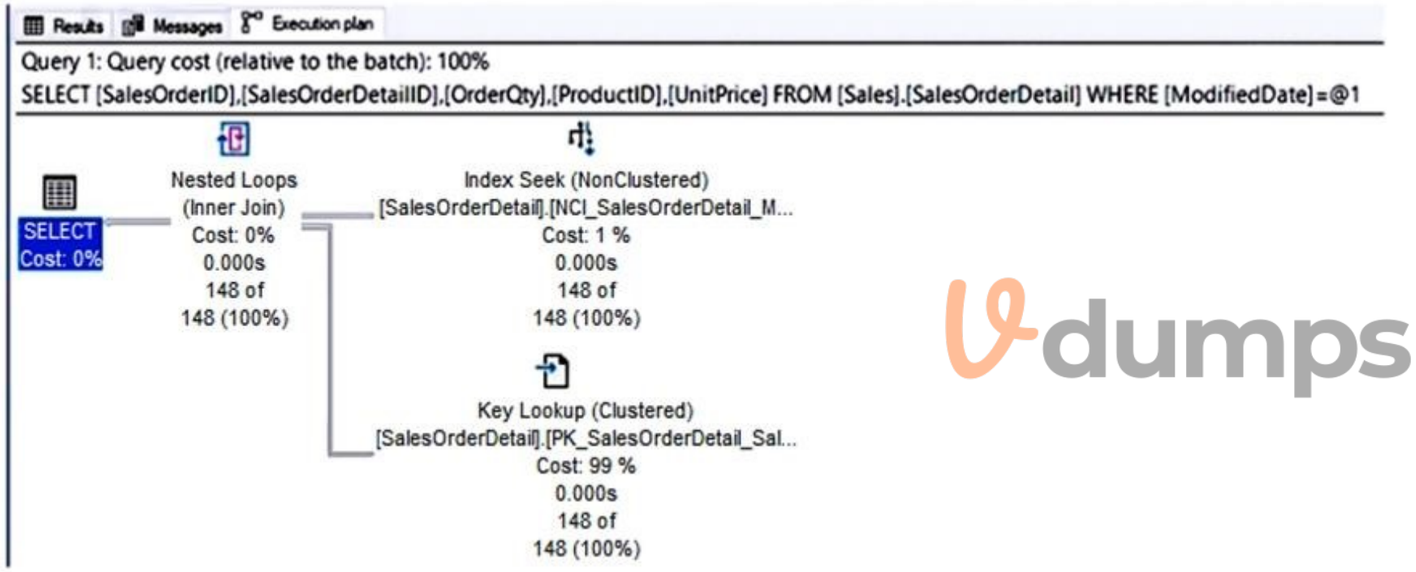

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

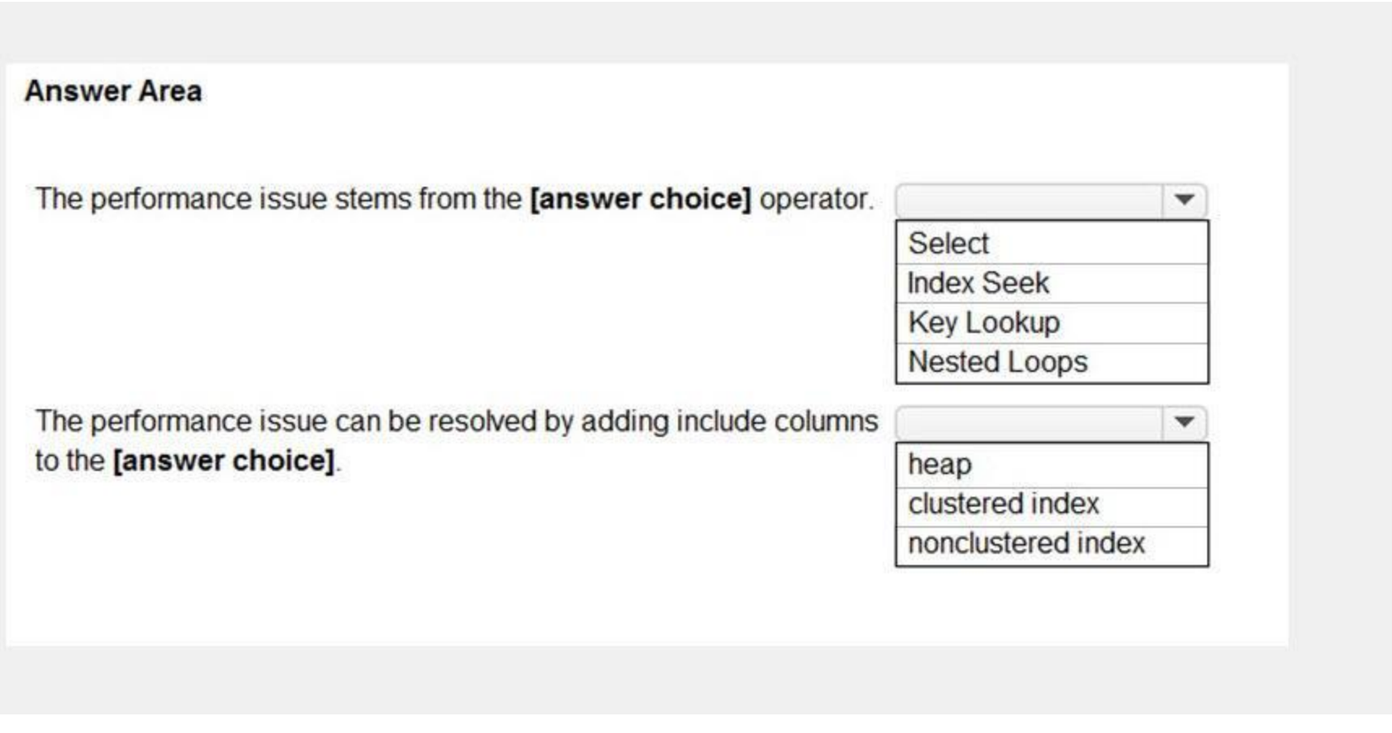

Answer Area:

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

Section:
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.

01 - Perform Automation of Tasks
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
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.

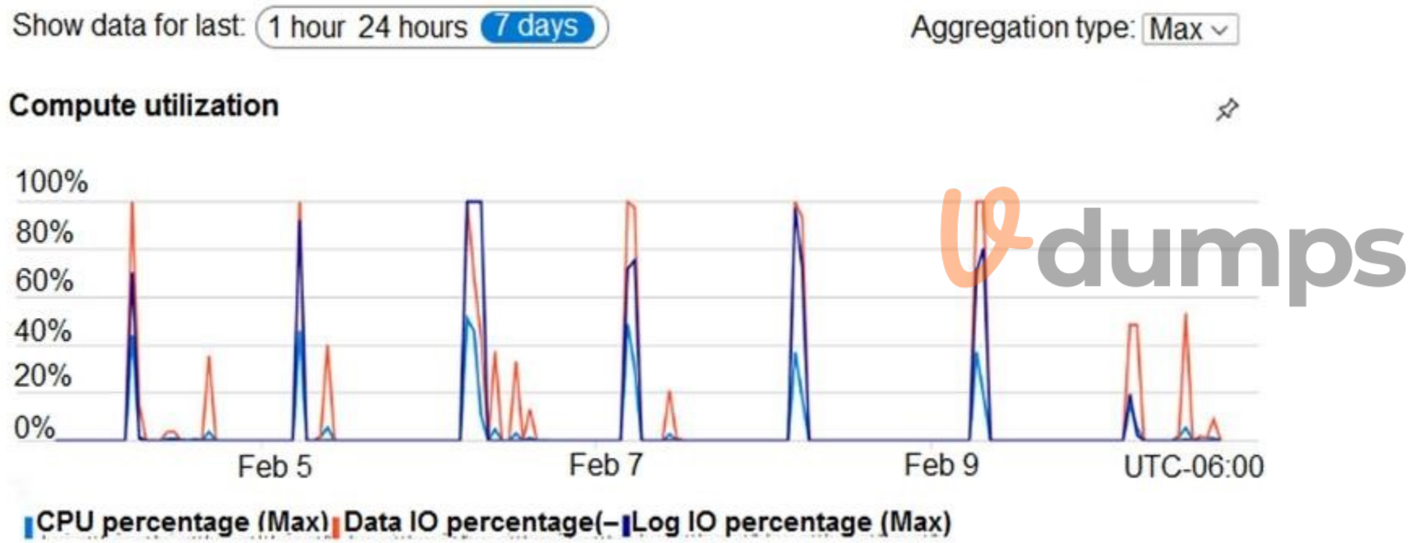

Role Assignments
For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

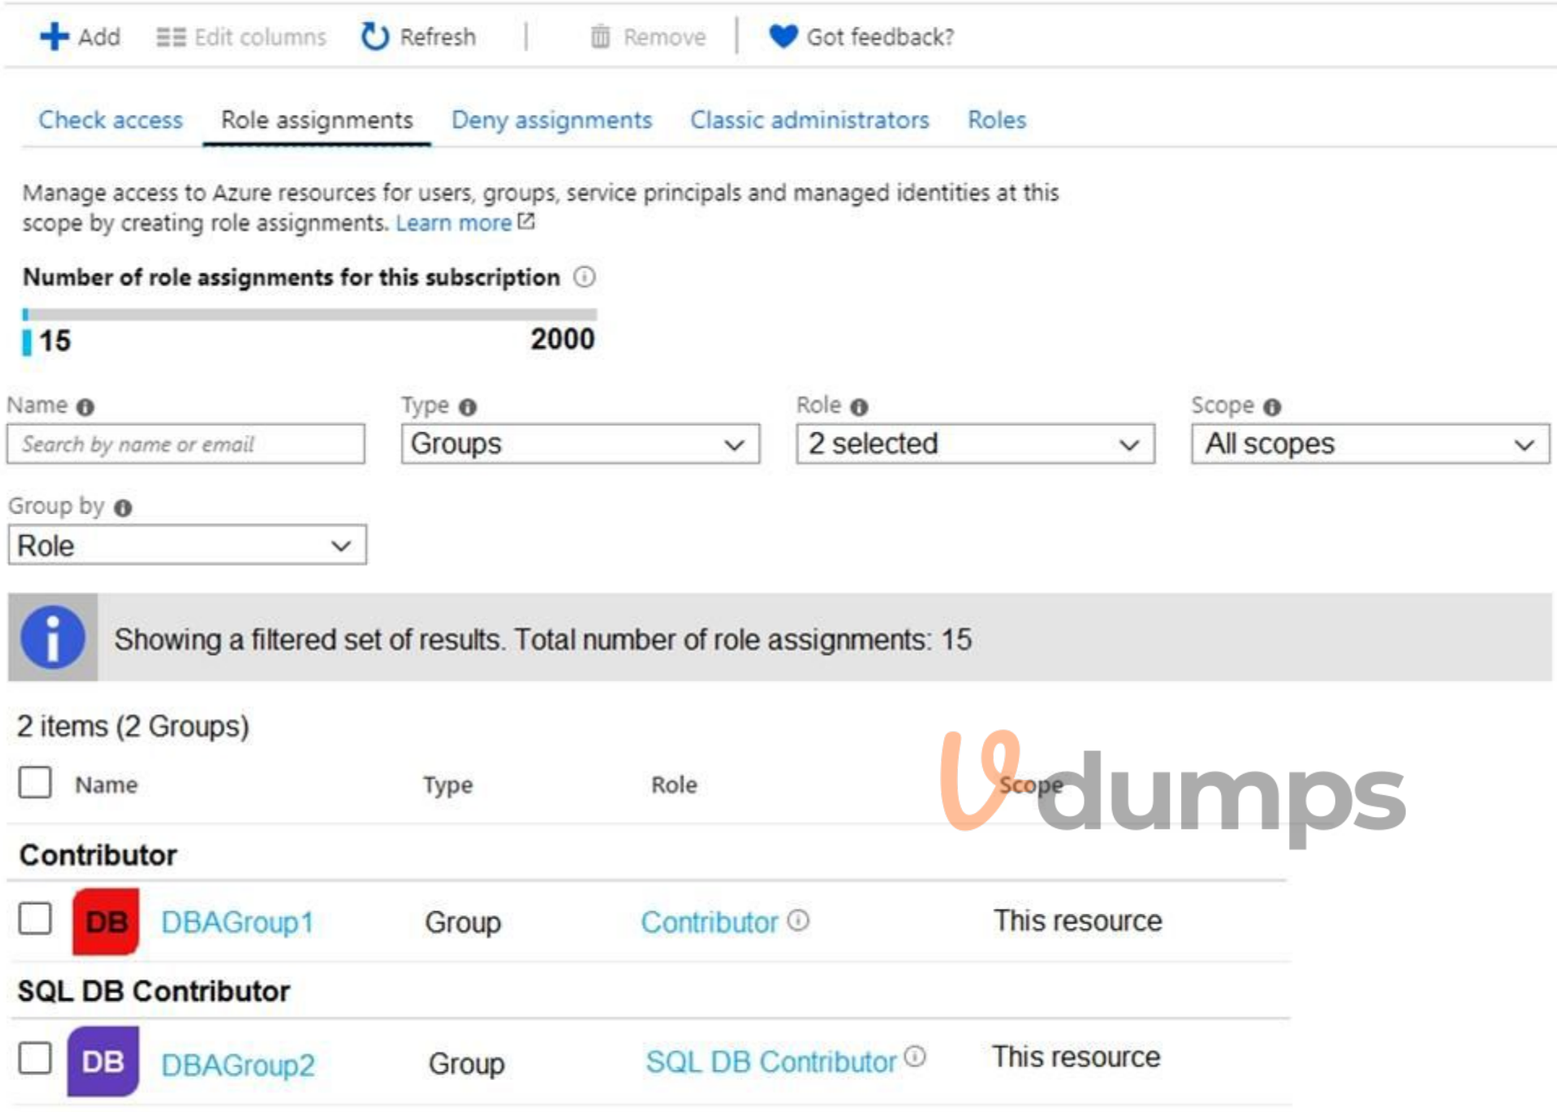

QUESTION 1
You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.
What should you do?

A. Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
B. Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
C. Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
D. Create an Azure Monitor alert rule that has a dynamic threshold and assign the alert rule to an action group.

Correct Answer: D
Section:
Explanation:

Reference:
https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/

02 - Perform Automation of Tasks

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

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

Requirements

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

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.

QUESTION 1
DRAG DROP
You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

Actions

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

Answer Area

Correct Answer:

Actions

- Create a SQL Server Agent job.
- Publish the runbook.
- Run sp_add_jobserver.

Answer Area

1. Create an Azure Automation account.
2. Import the SqlServer module.
3. Create a runbook that runs a PowerShell script.
4. Create and configure a schedule.

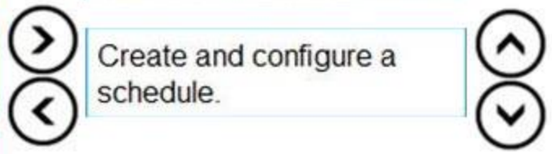

Section:

Explanation:

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

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

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

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

Steps:

1. Click on "runbooks" at the left panel and then click "add a runbook"
2. Choose "create a new runbook" and then give it a name and choose "Powershell" as the type of the runbook and then click on "create"

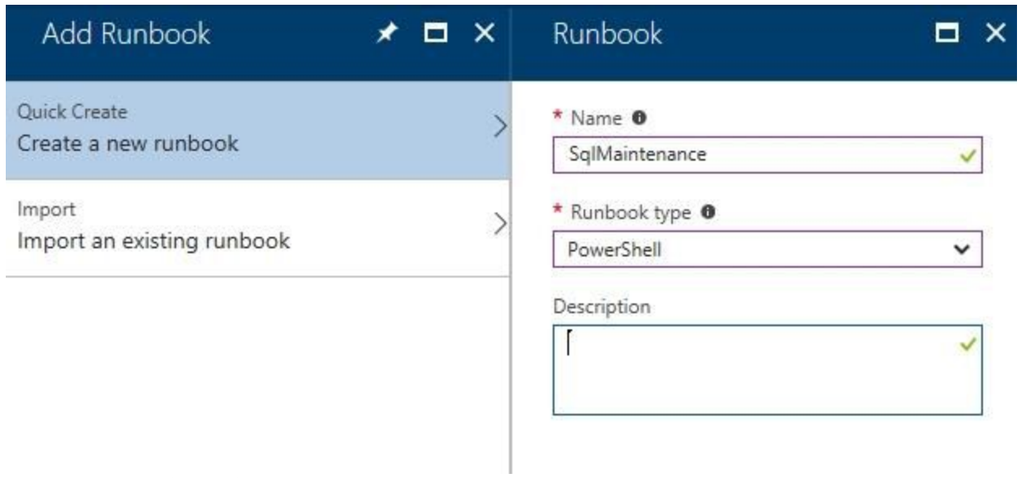

5. Schedule task (Step 4)
Steps:
1. Click on Schedules
2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.
Reference:
https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statistics-maintenance-using/ba-p/368974

03 - Perform Automation of Tasks

QUESTION 1
HOTSPOT
You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.
ADF1 contains the following pipelines:
P1:Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account P2:Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2
You need to configure P1 and P2 to maximize parallelism and performance.
Which dataset settings should you configure for the copy activity of each pipeline? To answer, select the appropriate options in the answer area.

Hot Area:

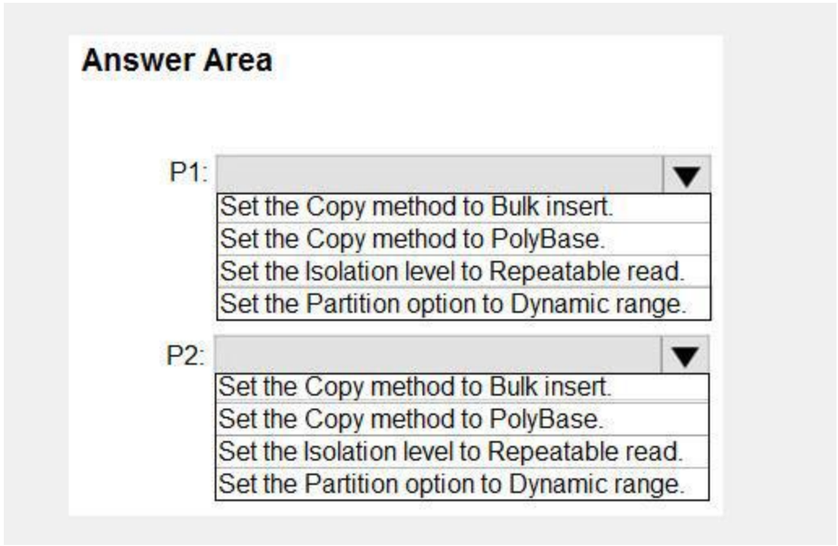

Answer Area:

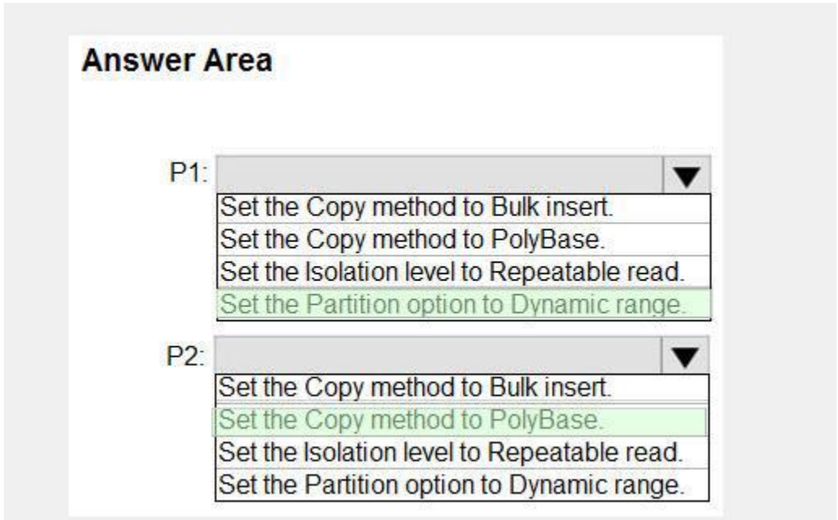

Section:
Explanation:
P1: Set the Partition option to Dynamic Range.
The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel.
P2: Set the Copy method to PolyBase
Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store. (Polybase supports Azure Blob storage and Azure Data Lake Store by default.)
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse
https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse

QUESTION 2

Hot Area:

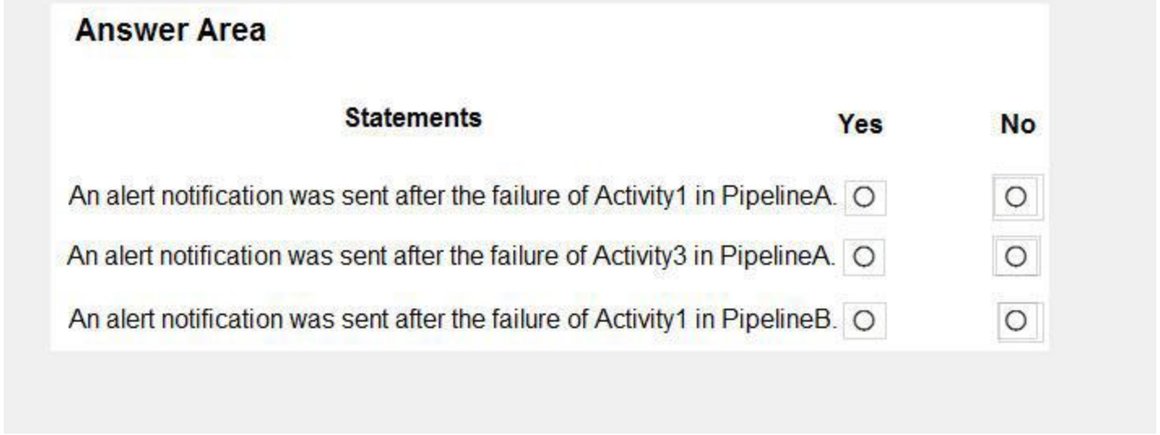

Answer Area:

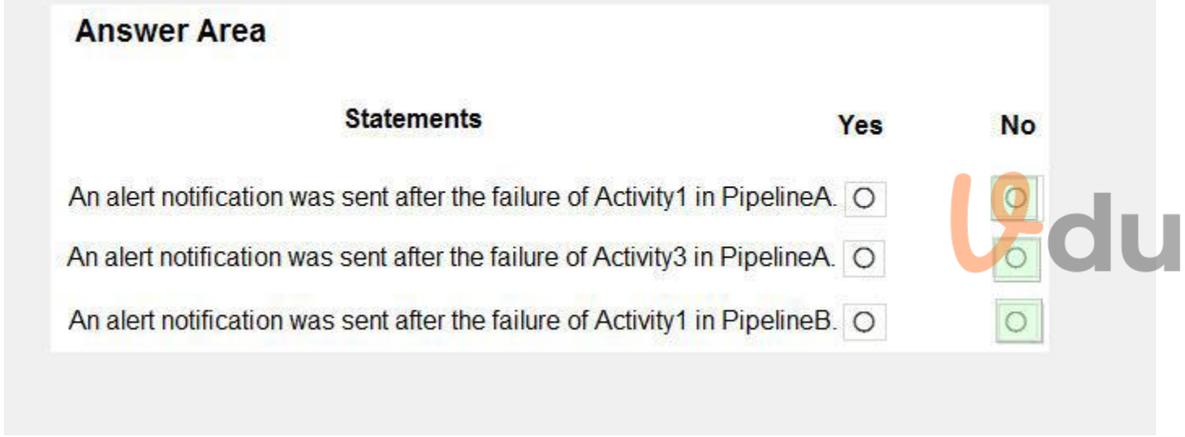

Section:
Explanation:
Box 1: No
Just one failure within the 5-minute interval.
Box 2: No
Just two failures within the 5-minute interval.
Box 3: No
Just two failures within the 5-minute interval.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview

QUESTION 3
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

Correct Answer: A
Section:
Explanation:
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/

QUESTION 4
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

Correct Answer: B
Section:
Explanation:
If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not a mapping flow,5 with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.
Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

QUESTION 5
Note: This question-is part of a series of questions that present the same scenario. Each question-in the series contains a unique solution that might meet the stated goals. Some question-sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Data Lake Storage account that contains a staging zone.
You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.
Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse. Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Must use an Azure Data Factory, not an Azure Databricks job.
Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

QUESTION 6
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 an Azure Databricks notebook, and then inserts the data into the data warehouse. Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not an Azure Databricks notebook, with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.
Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

QUESTION 7
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

Correct Answer: A
Section:
Explanation:
If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.
Reference:
https://docs.microsoft.com/en-US/azure/data-factory/transform-data

QUESTION 8
DRAG DROP
You have an Azure subscription that contains an Azure SQL managed instance named SQLMi1 and a SQL Agent job named Backupdb. Backupdb performs a daily backup of the databases hosted on SQLMi1.
You need to be notified by email if the job fails.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

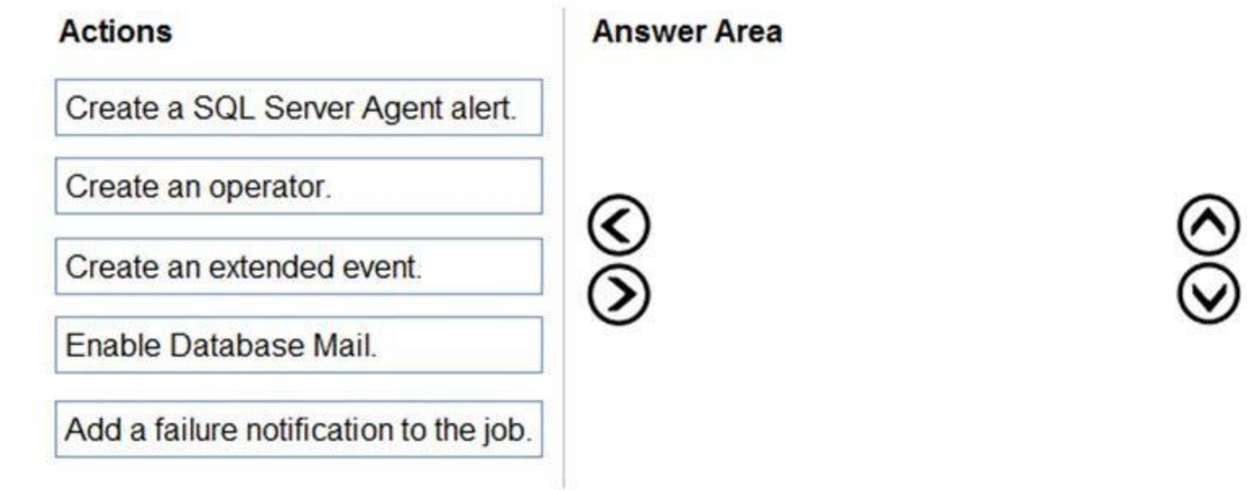

Correct Answer:

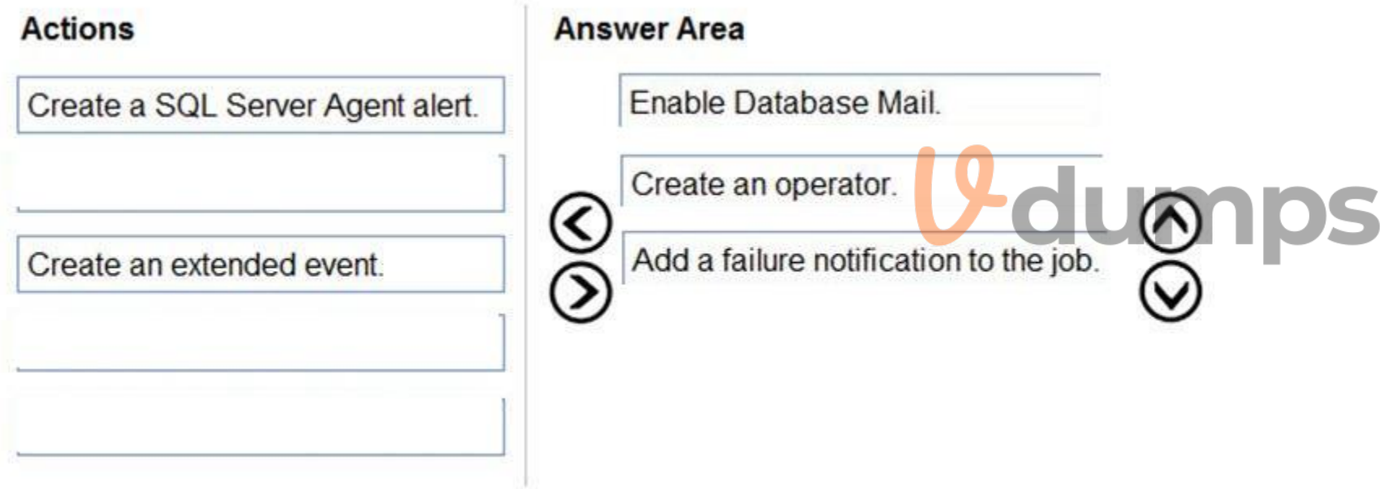

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance

QUESTION 9
DRAG DROP
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.

Select and Place:

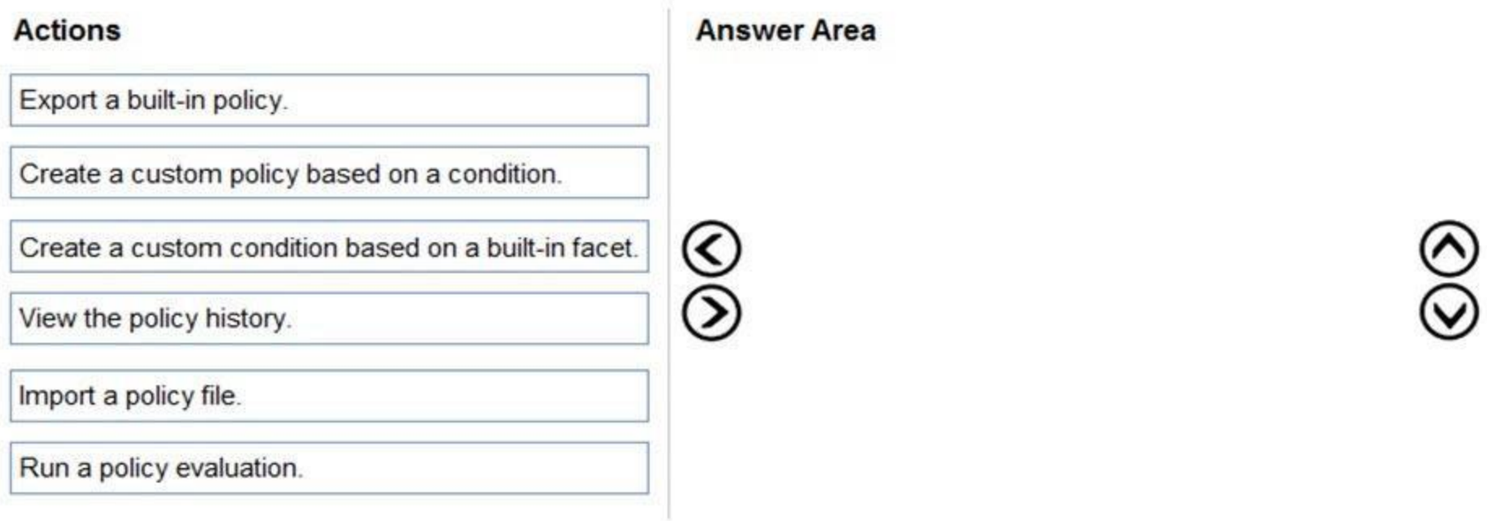

Correct Answer:

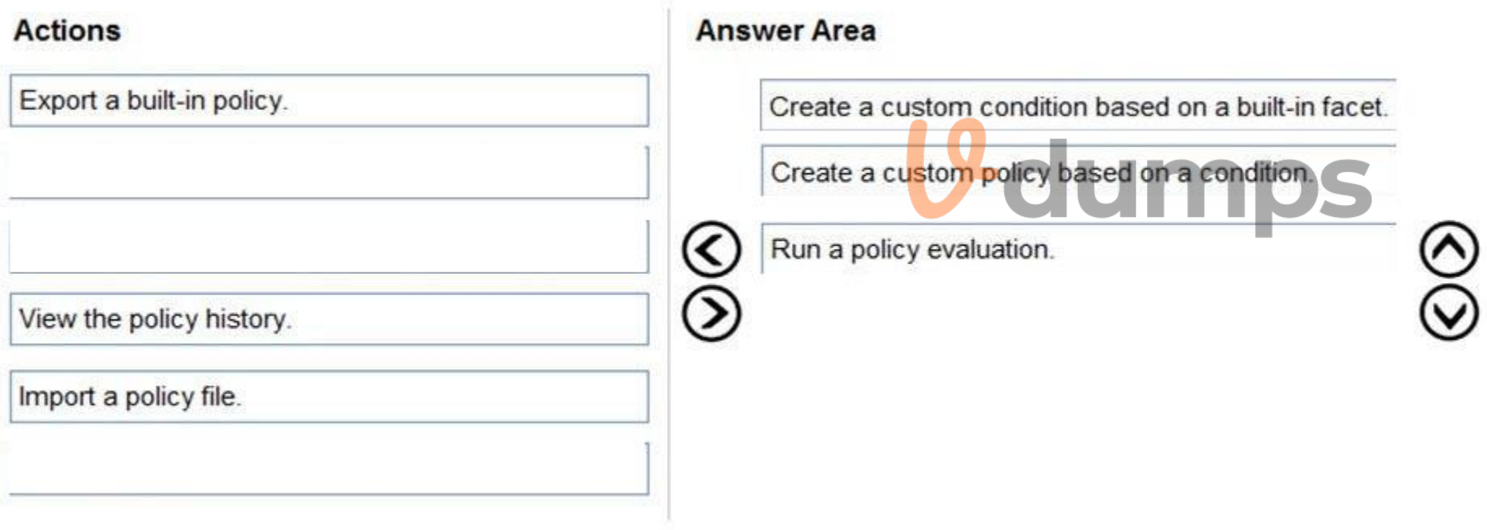

Section:
Explanation:
Reference:
https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/

QUESTION 10
You have an Azure SQL managed instance named SQLMI1 that hosts 10 databases.
You need to implement alerts by using Azure Monitor. The solution must meet the following requirements:
Minimize costs.
Aggregate Intelligent Insights telemetry from each database.
What should you do?

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

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

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry

QUESTION 11
You have an Azure SQL managed instance that hosts multiple databases.
You need to configure alerts for each database based on the diagnostics telemetry of the database.
What should you use?

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

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/metrics-diagnostic-telemetry-logging-streaming-export-configure?tabs=azure-portal#configure-the-streaming-export-of-diagnostic-telemetry

QUESTION 12
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

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

QUESTION 13
DRAG DROP
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.

Select and Place:

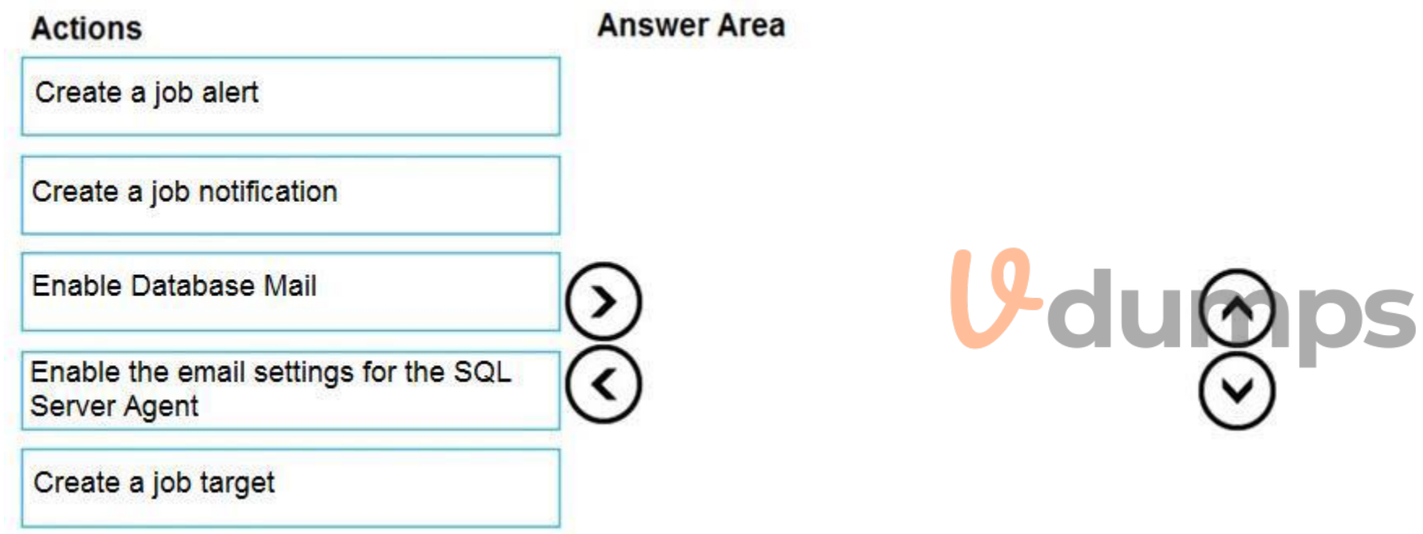

Correct Answer:

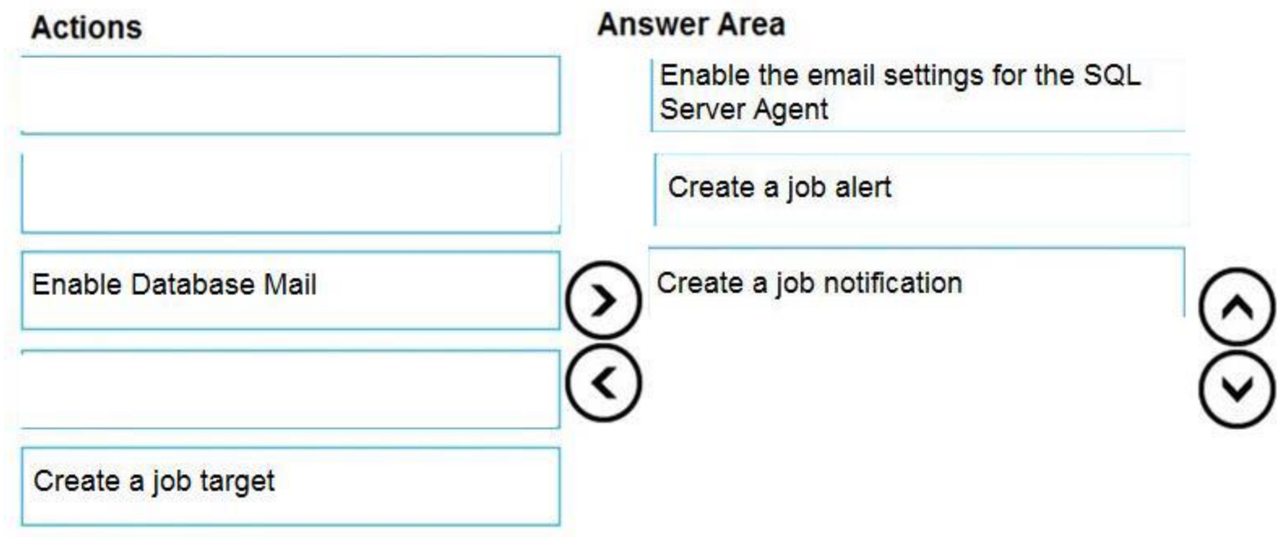

Section:
Explanation:
Step 1: Enable the email settings for the SQL Server Agent.
To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.
Step 2: Create a job alert
Step 3: Create a job notification
Example:

```
-- 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

QUESTION 14
DRAG DROP
You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

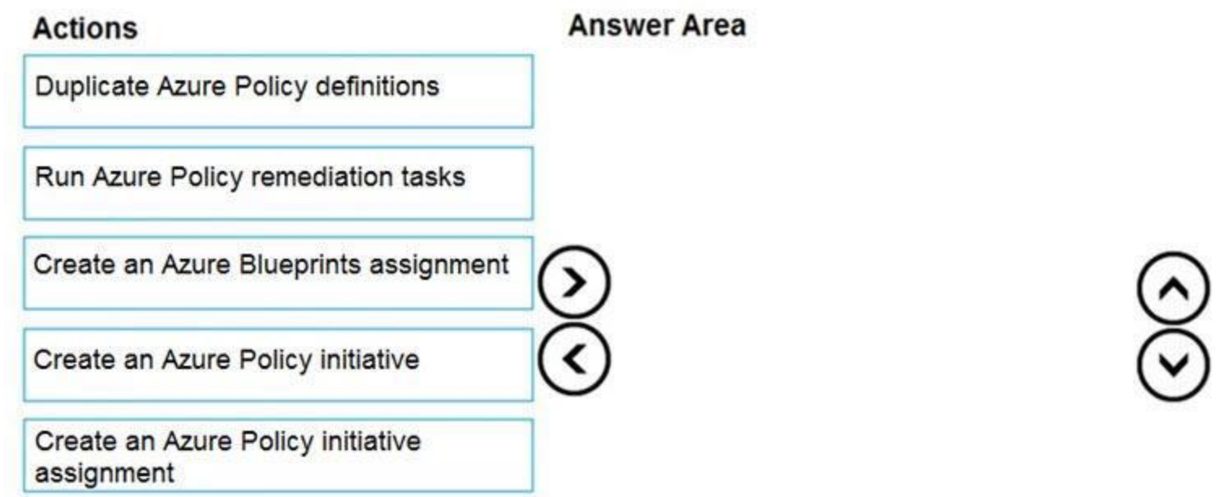

Correct Answer:

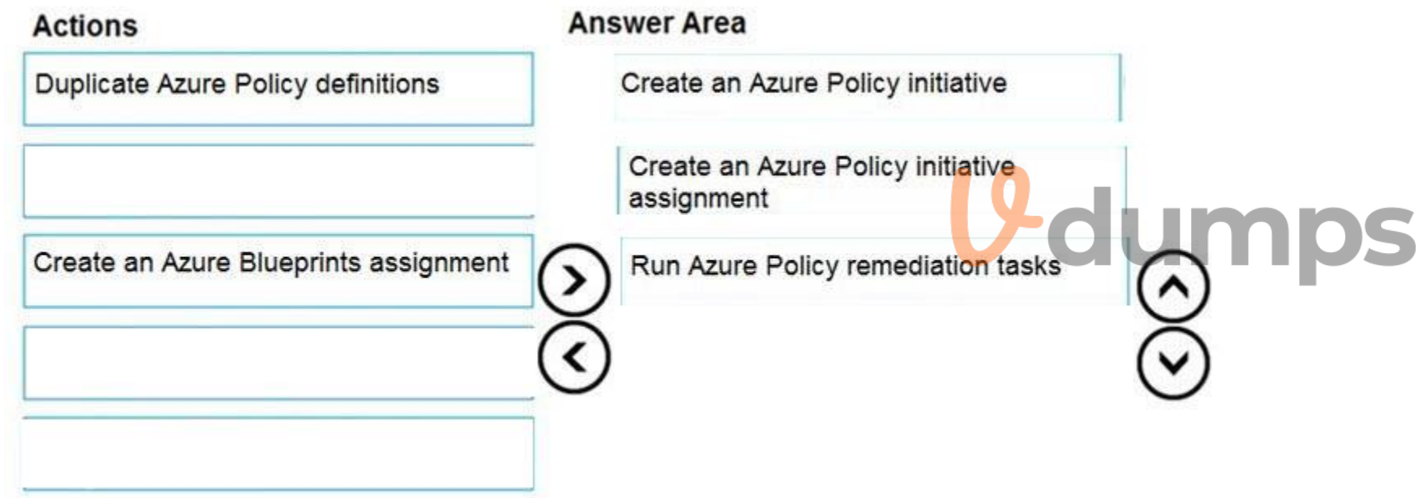

Section:
Explanation:
Step 1: Create an Azure Policy Initiative
The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.
With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.
Step 2: Create an Azure Policy Initiative assignment
Assign the initiative definition you created in the previous step.
Step 3: Run Azure Policy remediation tasks
To apply the Policy Initiative to the existing SQL databases.
Reference:
https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage

QUESTION 15
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

Correct Answer: B
Section:
Explanation:
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

QUESTION 16
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.

Correct Answer: D
Section:
Explanation:
Reference: https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/

QUESTION 17
You have an Azure Data Factory pipeline that performs an incremental load of source data to an Azure Data Lake Storage Gen2 account.
Data to be loaded is identified by a column named LastUpdatedDate in the source table.
You plan to execute the pipeline every four hours.
You need to ensure that the pipeline execution meets the following requirements:
Automatically retries the execution when the pipeline run fails due to concurrency or throttling limits. Supports backfilling existing data in the table.
Which type of trigger should you use?

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

Correct Answer: A
Section:

Explanation:
The Tumbling window trigger supports backfill scenarios. Pipeline runs can be scheduled for windows in the past. Incorrect Answers:
D: Schedule trigger does not support backfill scenarios. Pipeline runs can be executed only on time periods from the current time and the future.
Reference: https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipeline-execution-triggers

01 - Plan and Implement a High Availability and Disaster Recovery (HADR) Environment
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.
When you are ready to answer a question, click the Question button to return to the question.
Overview
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.

QUESTION 1
HOTSPOT
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.

Hot Area:

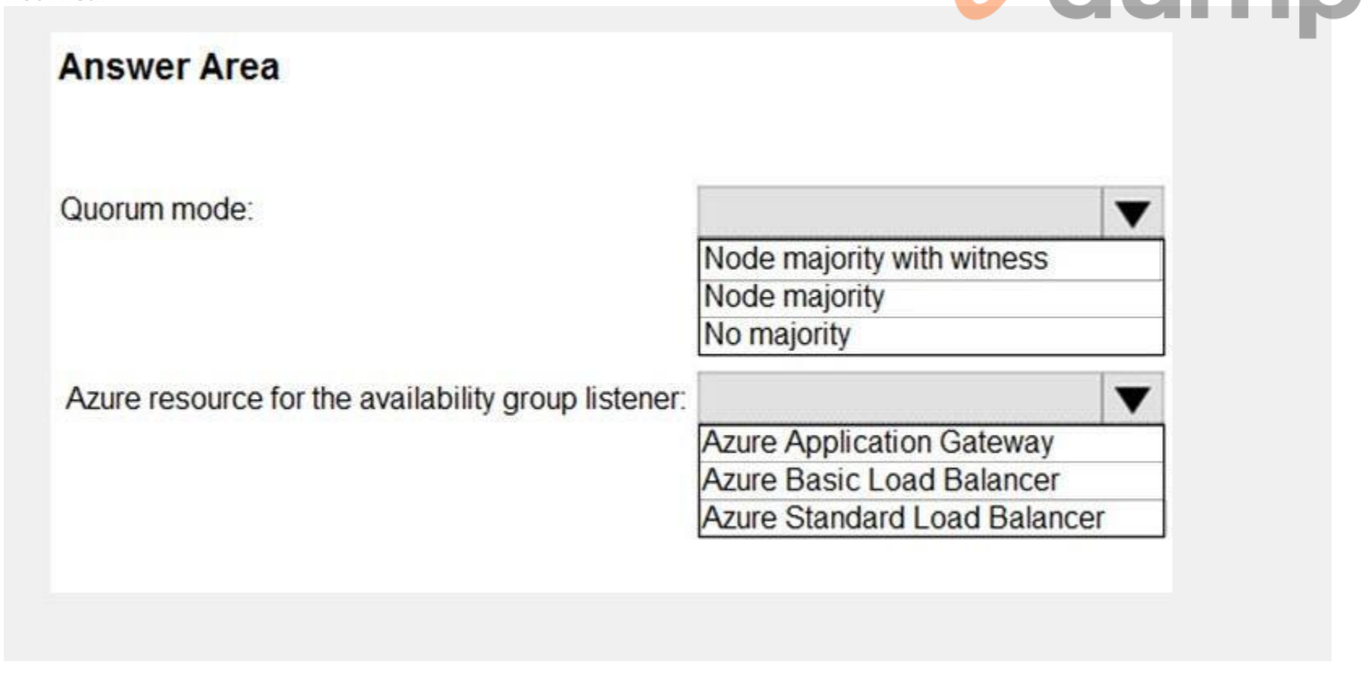

Answer Area:

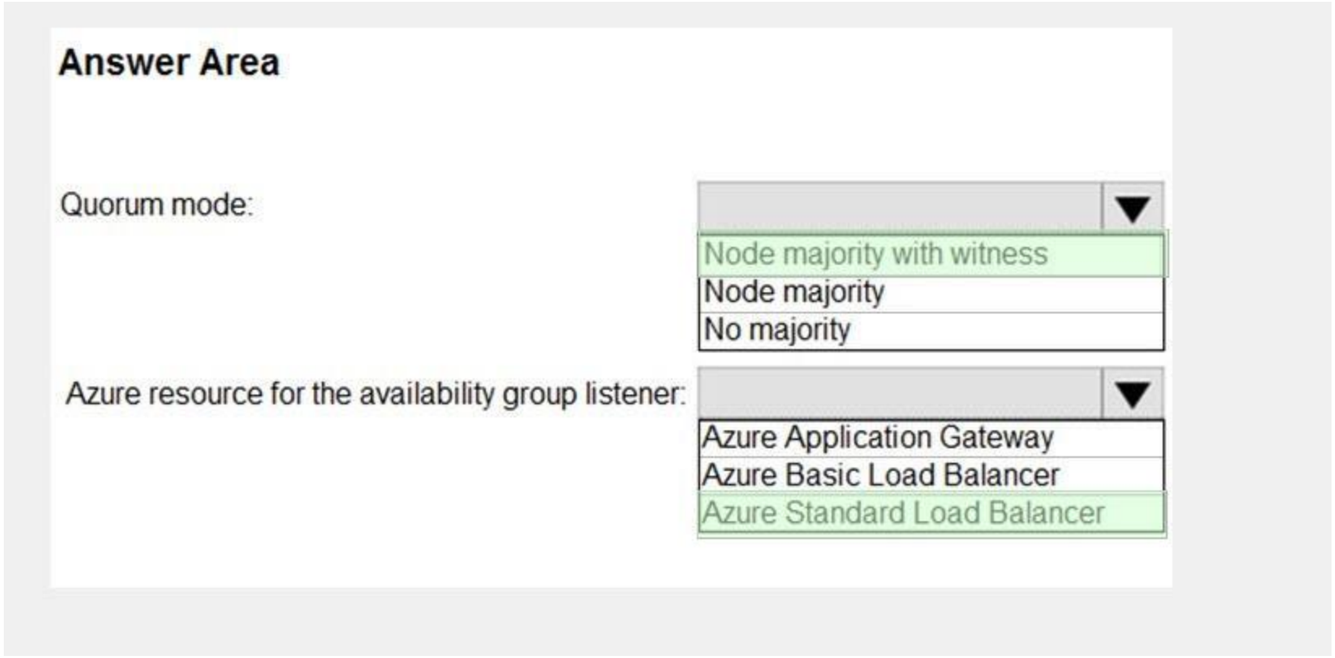

Section:
Explanation:
Box 1: Node majority with witness
As a general rule when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster contains an even number of voting nodes, you should configure a disk witness or a file share witness.
Note: Mode: Node majority with witness (disk or file share)
Nodes have votes. In addition, a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated disk witness or a designated file share witness.
Box 2: Azure Standard 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.
Scenario: Business Requirements
Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.
Incorrect Aswers:
Basic Balancer: No SLA is provided for Basic Load Balancer.
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/

QUESTION 2
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.

Correct Answer: B

Section:
Explanation:

QUESTION 3
You need to provide an implementation plan to configure data retention for ResearchDB1. The solution must meet the security and compliance requirements.
What should you include in the plan?

A. Configure the Deleted databases settings for ResearchSrv01.
B. Deploy and configure an Azure Backup server.
C. Configure the Advanced Data Security settings for ResearchDB1.
D. Configure the Manage Backups settings for ResearchSrv01.

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/long-term-backup-retention-configure

02 - Plan and Implement a High Availability and Disaster Recovery (HADR) Environment
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.
When you are ready to answer a question, click the Question button to return to the question.
Overview
Existing Environment
Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.
Active Directory
Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.
Database Environment
Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

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

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.
The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.
Current Business Model
Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers. The customers receive the files by using FTP.
Requirements
Planned Changes
Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an

Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.
The estimated size of each PaaS database is 1 TB.
Contoso plans to implement the following changes:
Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.
Start onboarding customers to the new PaaS solution within six months.
Business Goals
Contoso identifies the following business requirements:
Use built-in Azure features whenever possible.
Minimize development effort whenever possible.
Minimize the compute costs of the PaaS solutions.
Provide all the customers with their own copy of the database by using the PaaS solution.
Provide the customers with different table and row access based on the customer's service agreement. In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover. Ensure that users of the PaaS solution can create their own database objects but 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.

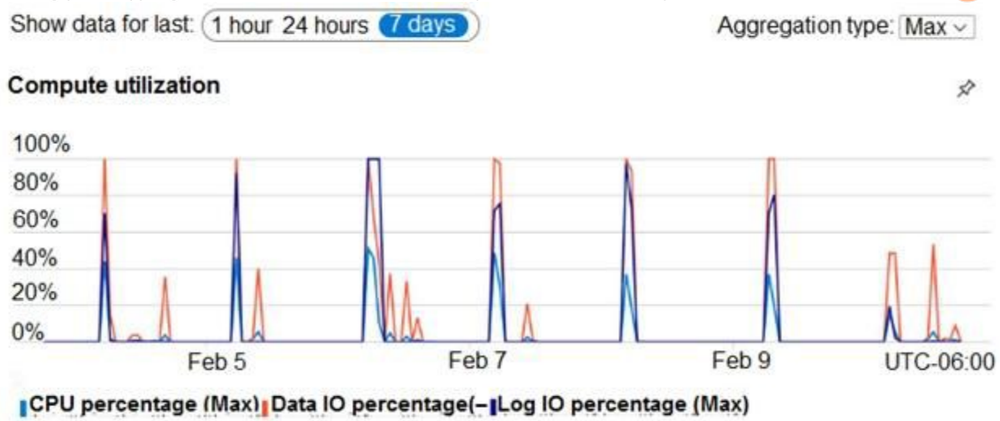

Role Assignments
For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

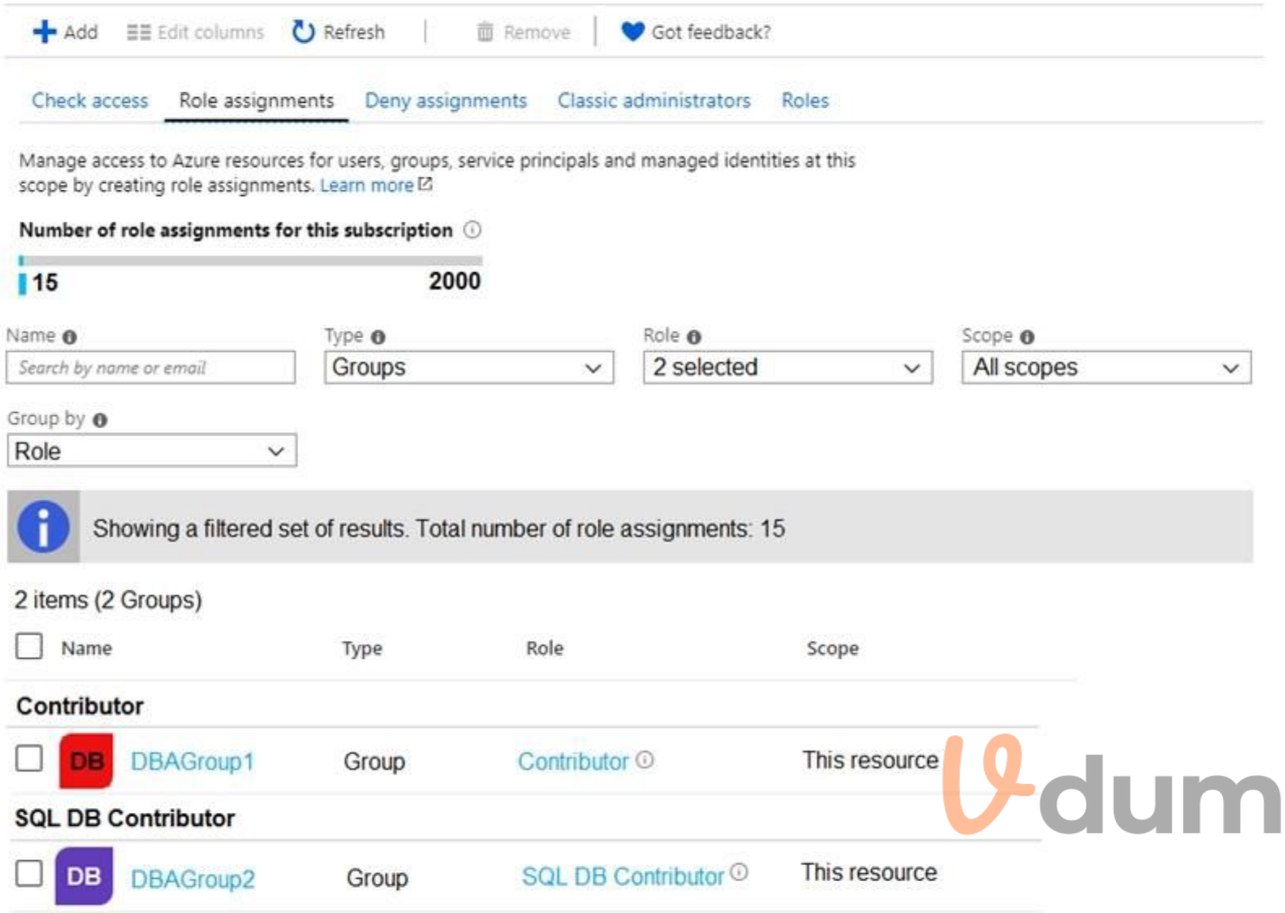

QUESTION 1
What should you implement to meet the disaster recovery requirements for the PaaS solution?

A. Availability Zones
B. failover groups
C. Always On availability groups
D. geo-replication

Correct Answer: B
Section:
Explanation:
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 georeplication 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

03 - Plan and Implement a High Availability and Disaster Recovery (HADR) Environment

QUESTION 1
You are planning disaster recovery for the failover group of an Azure SQL Database managed instance. Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs. You set the Read/Write failover policy to Automatic.
What are two results of the configuration? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Correct Answer: A, E
Section:
Explanation:
A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region. E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.
Incorrect Answers:
C: individual SQL Managed Instance databases cannot be added to or removed from a failover group.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview

QUESTION 2
You have an Azure SQL database named DB1.
You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs. Which deployment option and pricing tier should you configure?

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

Correct Answer: C
Section:
Explanation:
By default, the cluster of nodes for the premium availability model is created in the same datacenter.With the introduction of Azure Availability Zones, SQL Database can place different replicas of theBusiness Critical database to different availability zones in the same region. To eliminate a singlepoint of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zoneredundant configuration in the Premium or Business Critical service tiers does not create additionaldatabase redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration,you can make your Premium or Business Critical databases resilient to a much larger set of failures,including catastrophic datacenter outages, without any changes to the application logic. You can alsoconvert any existing Premium or Business Critical databases or pools to the zone redundantconfiguration.

QUESTION 3
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:
During normal operations, provide at least two readable copies of Sales.
Ensure that Sales remains available if a datacenter fails.
Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

QUESTION 4
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure SQL database named Sales.
You need to implement disaster recovery for Sales to meet the following requirements:
During normal operations, provide at least two readable copies of Sales.
Ensure that Sales remains available if a datacenter fails.
Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.
Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
Explanation:
Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

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

You need to implement disaster recovery for Sales to meet the following requirements:
During normal operations, provide at least two readable copies of Sales.
Ensure that Sales remains available if a datacenter fails.
Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones. Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.
By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla

QUESTION 6
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.
You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.
Solution: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.
Does this meet the goal?

A. Yes
B. No

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

QUESTION 7
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

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

QUESTION 8
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

Correct Answer: A
Section:
Explanation:
The REPLACE option overrides several important safety checks that restore normally performs. The overridden checks are as follows: Restoring over an existing database with a backup taken of another database. With the REPLACE option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.
Reference: https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

QUESTION 9
You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica. You need to initiate a full backup of DB1 on SQL2.
Which statement should you run?

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

Correct Answer: B
Section:
Explanation:
BACKUP DATABASE supports only copy-only full backups of databases, files, or filegroups when it's executed on secondary replicas. Copy-only backups don't impact the log chain or clear the differential bitmap. Incorrect Answers:
A: Differential backups are not supported on secondary replicas. The software displays this error because the secondary replicas support copy-only database backups.
Reference: https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups

QUESTION 10
You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations. What should you use?

A. An Azure SQL Database elastic pool
B. SQL Server on Azure virtual machines
C. an Azure SQL Database managed instance
D. Azure SQL databases

Correct Answer: A
Section:
Explanation:
Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price.
Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

QUESTION 11
You have 10 Azure virtual machines that have SQL Server installed.
You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups. What should you include in the backup strategy?

A. Automated Backup in the SQL virtual machine settings
B. Azure Backup
C. Azure Site Recovery
D. SQL Server Agent jobs

Correct Answer: B
Section:
Explanation:
Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault. There are several advantages that this solution provides, especially for Enterprises.
Reference: https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/backup-restore#azbackup

QUESTION 12
You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:
Support failovers that do not require client applications to change their connection strings.
Replicate the database to a secondary Azure region.
Support failover to the secondary region.
What should you include in the recommendation?

A. failover groups
B. transactional replication
C. Availability Zones
D. geo-replication

Correct Answer: A
Section:
Explanation:

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region). Incorrect Answers:
C: Availability Zones are unique physical locations within a region. Each zone is made up of one or more datacenters equipped with independent power, cooling, and networking.
Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview

QUESTION 13
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

Correct Answer: B
Section:
Explanation:
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

QUESTION 14
You have SQL Server on Azure virtual machines in an availability group.
You have a database named DB1 that is NOT in the availability group.
You create a full database backup of DB1.
You need to add DB1 to the availability group.
Which restore option should you use on the secondary replica?

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

Correct Answer: B
Section:
Explanation:
Prepare a secondary database for an Always On availability group requires two steps:
1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY
2. Join the restored database to the availability group.
Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server

QUESTION 15

You plan to move two 100-GB databases to Azure.
You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations. What should you use?

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

Correct Answer: A
Section:

QUESTION 16
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

Correct Answer: B
Section:
Explanation:
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

QUESTION 17
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

Correct Answer: B
Section:
Explanation:
Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option. Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

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

QUESTION 18
DRAG DROP
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, and transaction log backups every five minutes.
The database fails on Wednesday at 10:00.
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.

Select and Place:

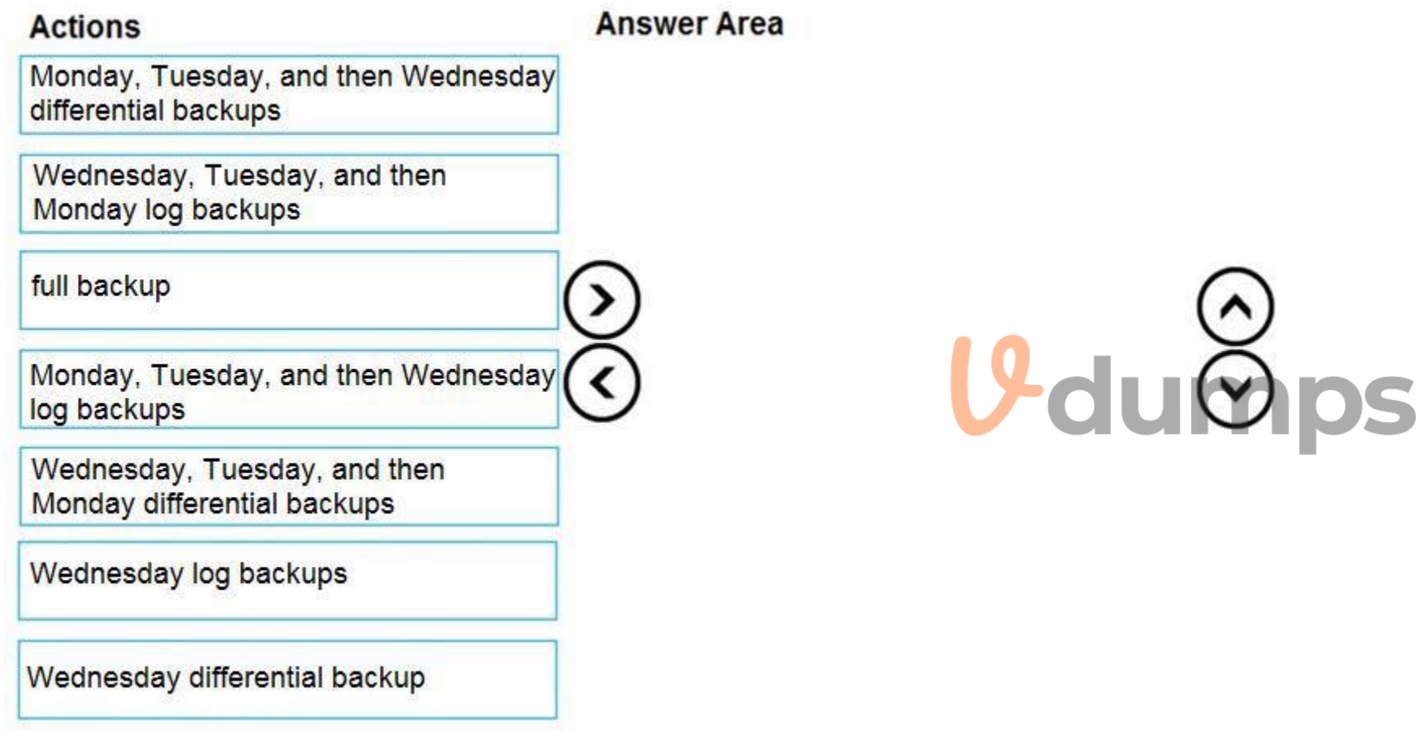

Correct Answer:

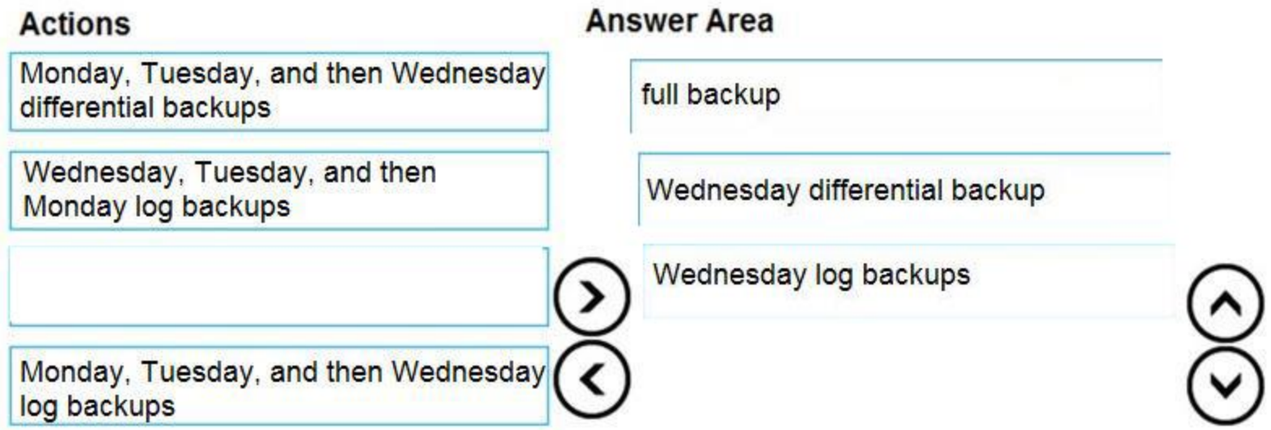

Section:
Explanation:

QUESTION 19
You have an Azure subscription that uses a domain named contoso.com.
You have two Azure VMs named DBServer1 and DBServer2. Each of them hosts a default SQL Server instance. DBServer1 is in the East US Azure region and contains a database named DatabaseA. DBServer2 is in the West US Azure region.
DBServer1 has a high volume of data changes and low latency requirements for data writes.
You need to configure a new availability group for DatabaseA. The secondary replica will reside on DBServer2.
What should you do?

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

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-ver15

QUESTION 20
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

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure

QUESTION 21
HOTSPOT
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.

Hot Area:

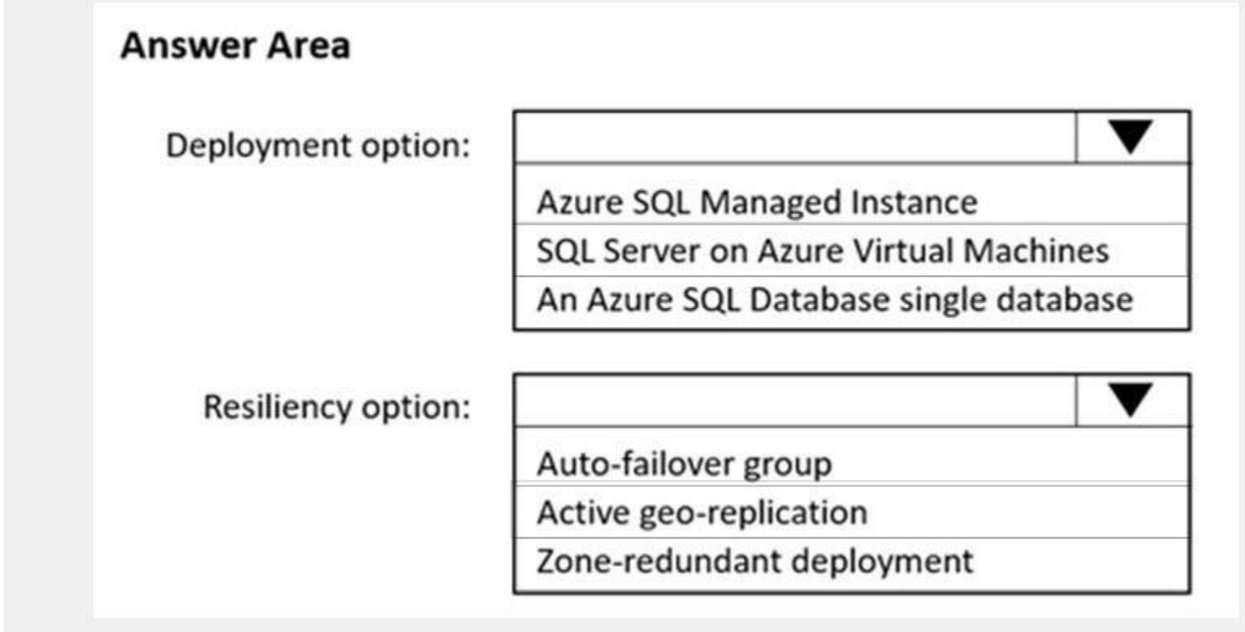

Answer Area:

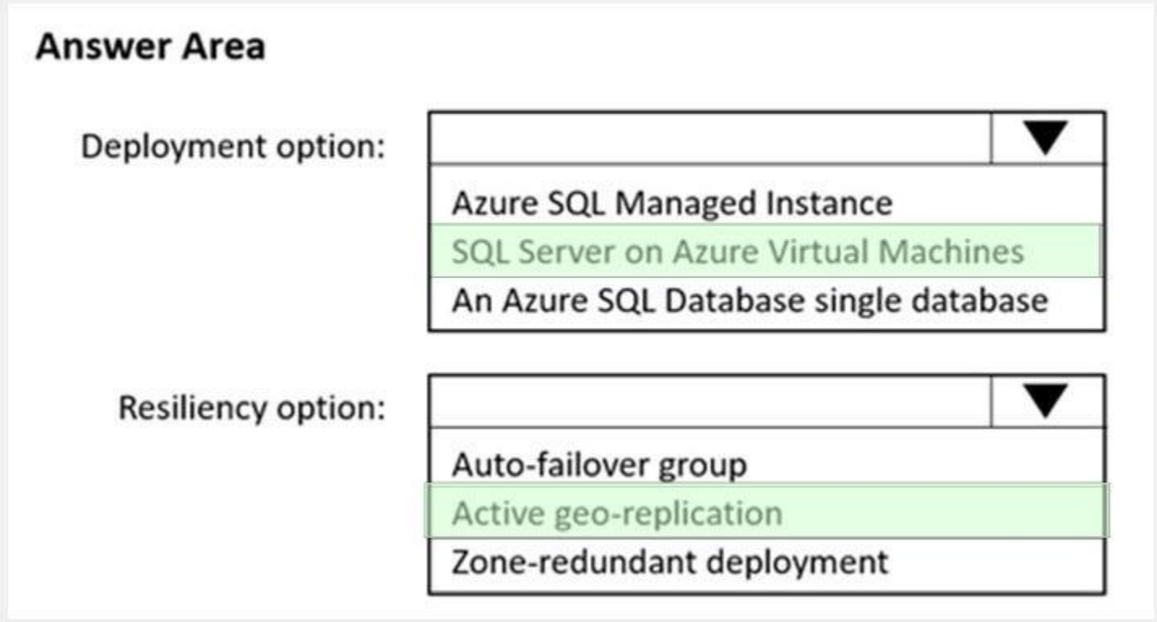

Section:
Explanation:
Box 1: SQL Server on Azure VMs
SQL Server on Azure Virtual Machines can take advantage of Automated Backup, which regularly creates backups of your database to blob storage. You can also manually use this technique.
Box 2: Active geo-replication
Geo-replication for services such as Azure SQL Database and Cosmos DB will create secondary replicas of your data across multiple regions. While both services will automatically replicate data within the same region, geo-replication protects you against a regional outage by enabling you to fail over to a secondary region.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-on-azure-vm-iaas-what-is-overview
https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/infrastructure-resiliency-azure

Exam L

QUESTION 1
HOTSPOT
You have a 50-TB Microsoft SQL Server database named DB1.
You need to reduce the time it takes to perform database consistency checks of DB1.
Which Transact-SQL command should you run? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

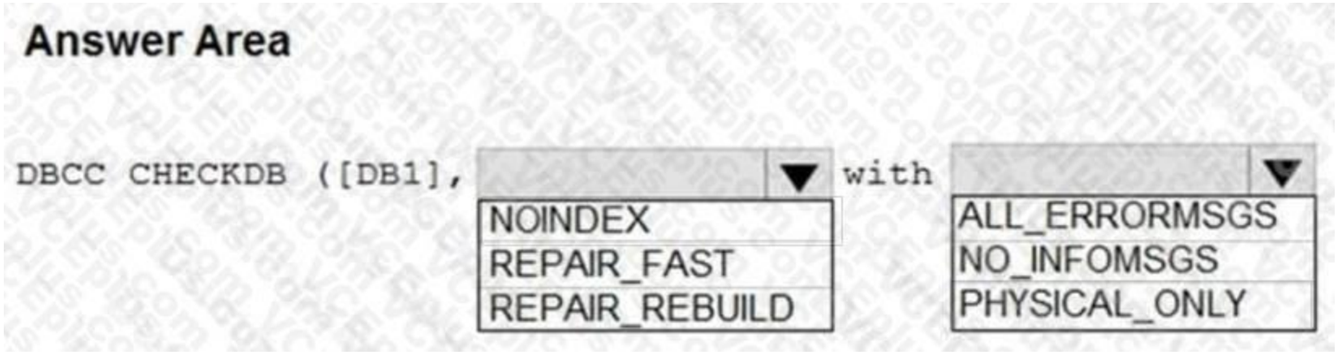

Answer Area:

Section:
Explanation:
Reference: https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-server-ver15

QUESTION 2
You have an instance of SQL Server on Azure Virtual Machine named SQL1.
You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort. Where should you store the metrics?

A. a Log Analytics workspace
B. Azure Event Hubs
C. Azure SQL Database
D. an Azure Blob storage container

Correct Answer: C
Section:

QUESTION 3
You have an Azure subscription that contains the resources shown in the following table.

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

App1 experiences transient connection errors and timeouts when it attempts to access db1 after extended periods of inactivity. You need to modify db1 to resolve the issues experienced by Appl as soon as possible, without considering immediate costs What you do?

A. Increase the number of vCores allocated to db1.
B. Decrease the auto-pause delay for db1.
C. Disable auto-pause delay for db1.
D. Enable automatic tuning for db1.

Correct Answer: D
Section:

QUESTION 4
HOTSPOT
You have a Microsoft SQL Server 2017 server that hosts five databases.
You Plan to migrate the databases to Azure.
You need to recommend a solution that meets the following requirements:
Automatically scales compute based on the workload demand
Provides per-second billing
What should you include in the recommendation? To answer, select the appropriate options in the answer area.

Hot Area:

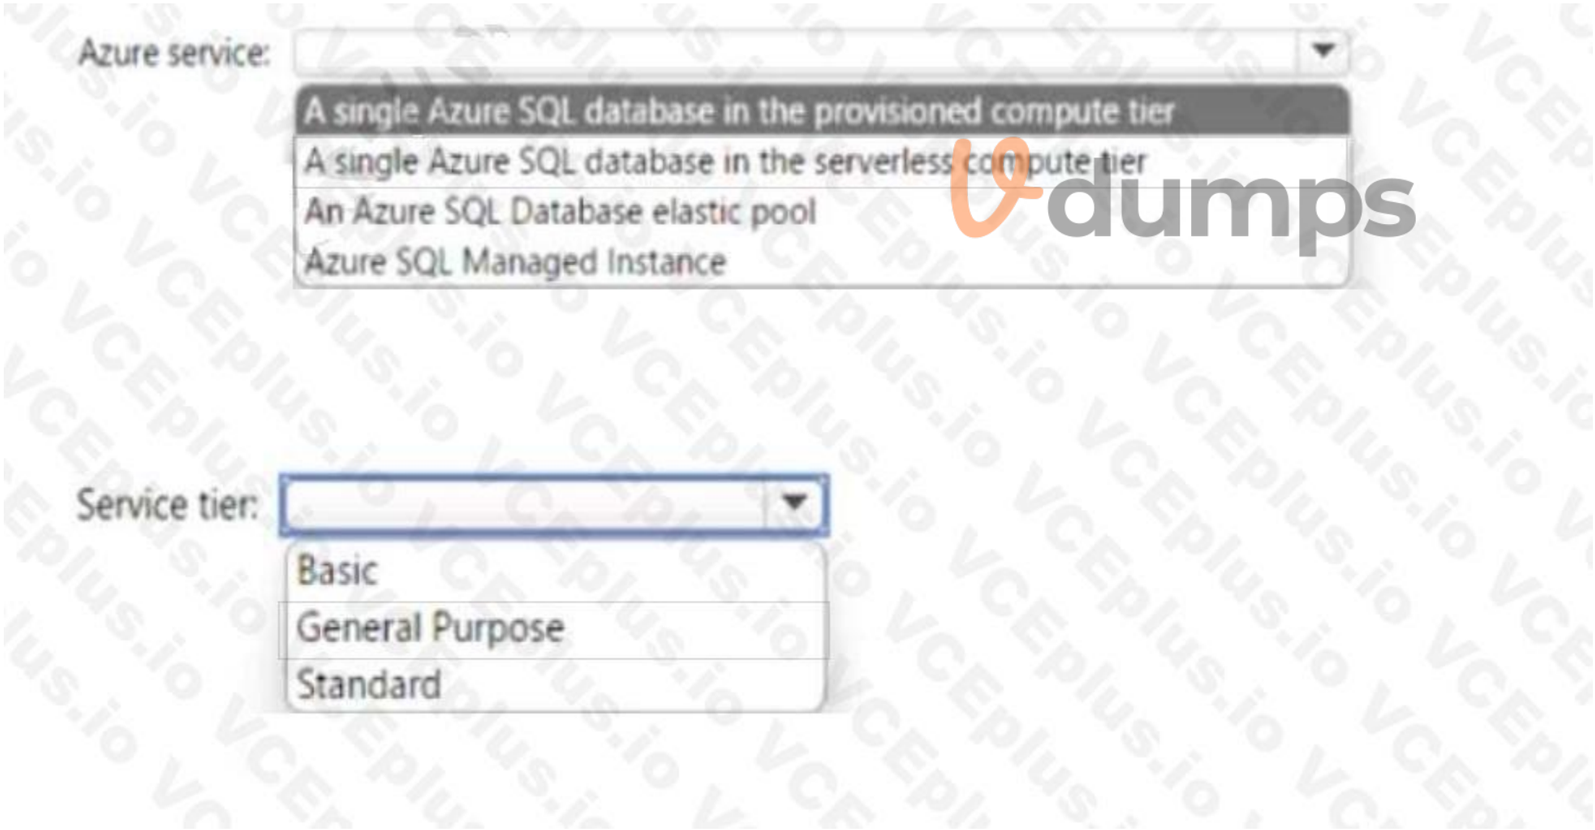

Answer Area:

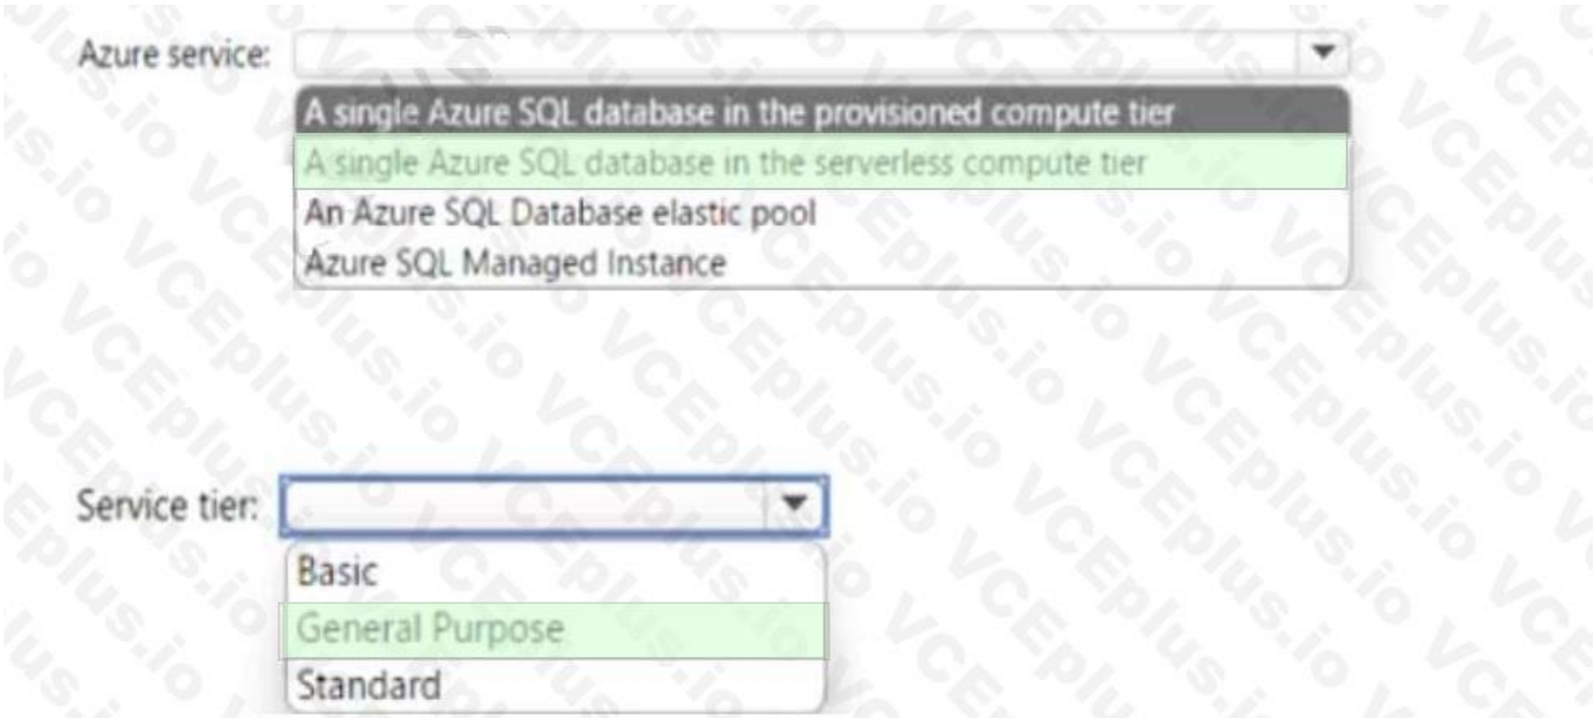

Section:
Explanation:

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

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

Correct Answer: B
Section:

QUESTION 6
DRAG DROP
You have a burstable Azure virtual machine named VMI that hosts an instance of Microsoft SQL Server. You need to attach an Azure ultra disk to VMI. The solution must minimize downtime on VMI.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

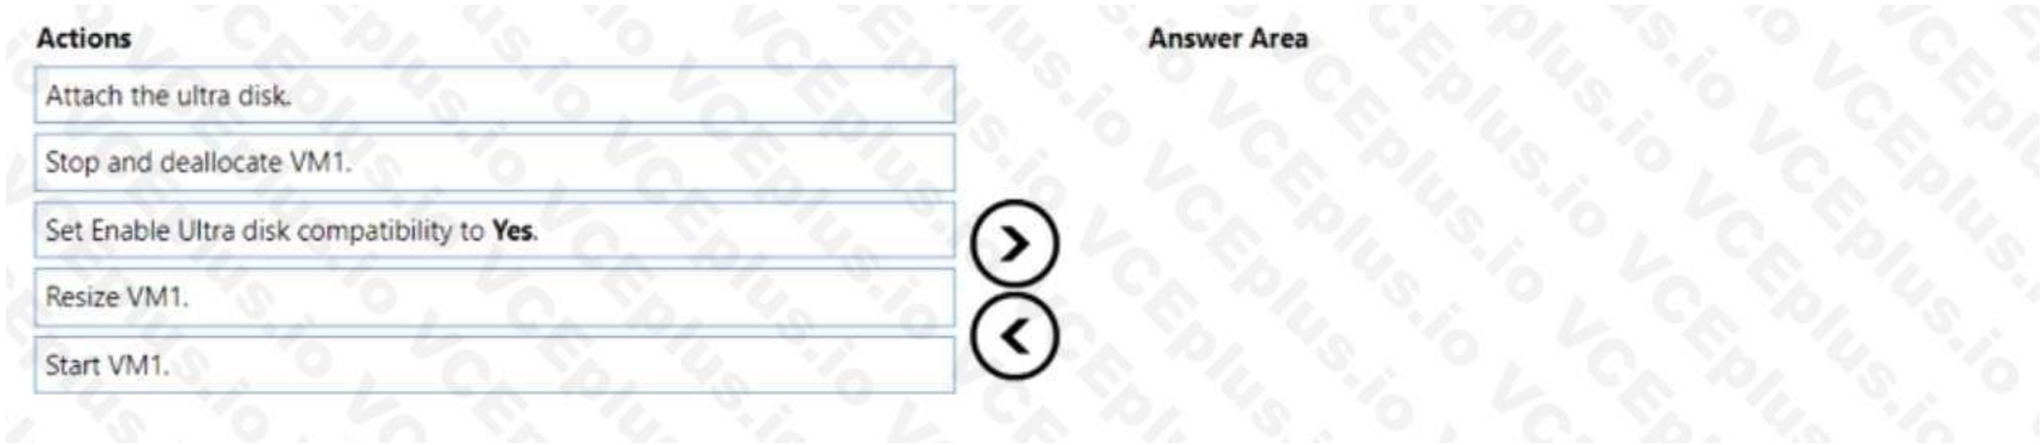

Correct Answer:

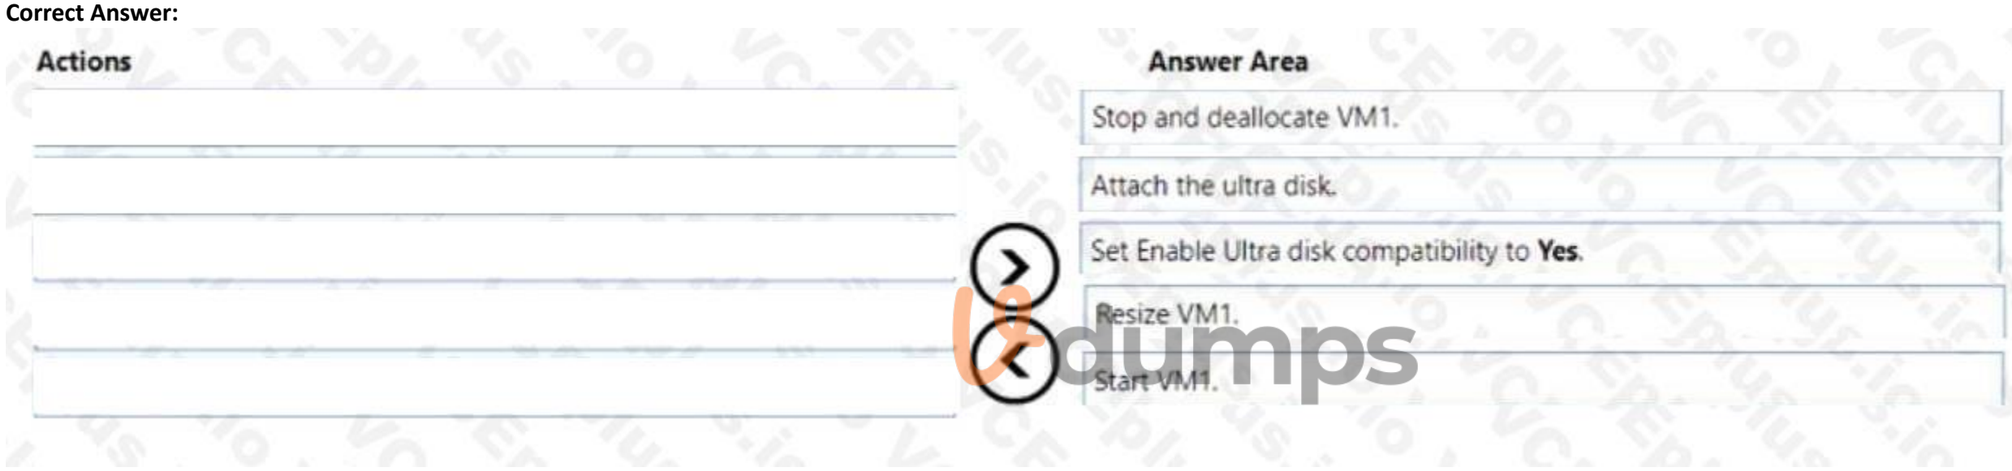

Section:
Explanation:

QUESTION 7
DRAG DROP
You have an Azure subscription.
You need to deploy an Azure SQL managed instance by using an Azure Resource Manager (ARM) template. The solution must meet the following requirements:
The SQL managed instance must be assigned a unique identity.
The SQL managed instance must be available in the event of an Azure datacenter outage.
How should you complete the template? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Select and Place:

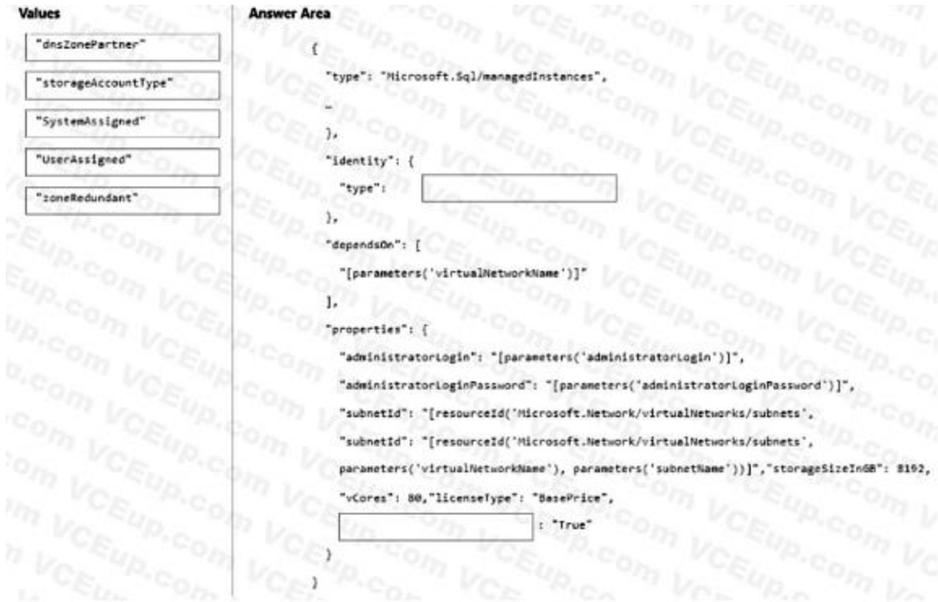

Correct Answer:

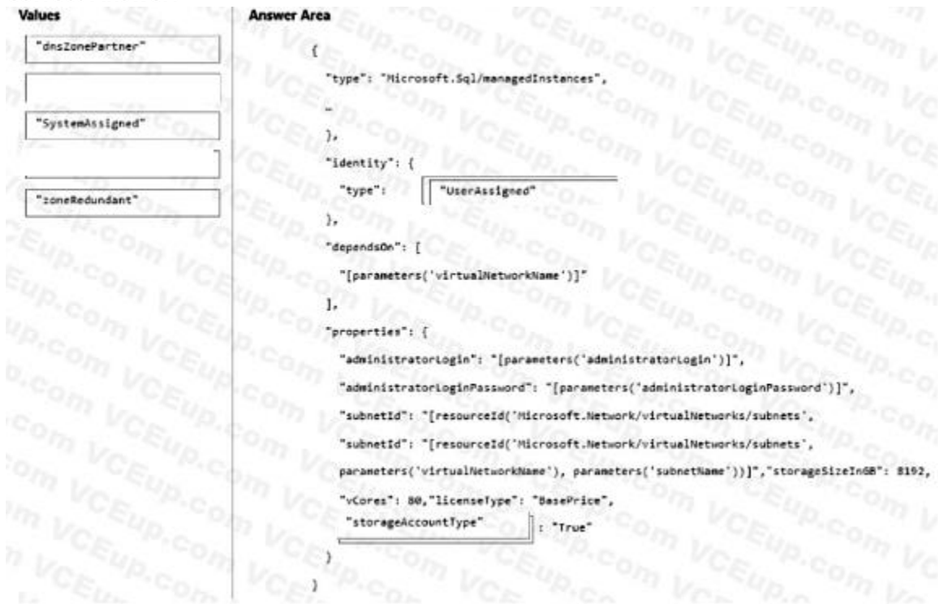

Section:
Explanation:

QUESTION 8
You have a database on a SQL Server on Azure Virtual Machines instance.
The current state of Query Store for the database is shown in the following exhibit.

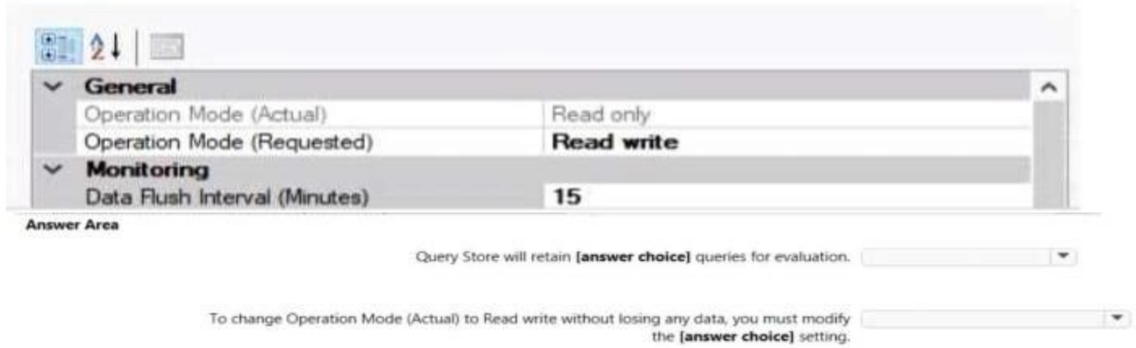

A.

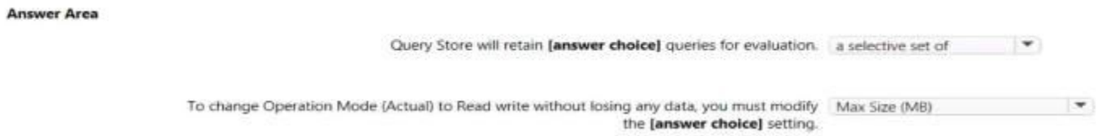

Correct Answer: A
Section:

QUESTION 9
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.

Correct Answer: B
Section:

QUESTION 10
You have an Azure subscription that contains a logical SQL server named Serverl. The master database of Serverl contains a user named Userl. You need to ensure that Userl can create databases on Serverl. Which database role should you assign to Userl?

A. db_owner
B. dbmanager
C. dbo
D. db_ddladmin

Correct Answer: B
Section:

QUESTION 11
You have an instance of SQL Server on Azure Virtual Machines that has a database named DB1.
You plan to implement Azure SQL Data Sync for DB1.

Which isolation level should you configure?

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

Correct Answer: B
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database

QUESTION 12
You have an Azure SQL managed instance.
You need to enable SQL Agent Job email notifications.
What should you do?

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

Correct Answer: C
Section:
Explanation:
Reference: https://techcommunity.microsoft.com/t5/azure-sql-blog/sending-emails-in-azure-sql-managed-instance/ba-p/386235

QUESTION 13
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

Correct Answer: A
Section:
Explanation:
Reference: https://www.sqlshack.com/automate-azure-sql-database-indexes-and-statistics-maintenance/

QUESTION 14
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.

Correct Answer: B, C
Section:
Explanation:

QUESTION 15
HOTSPOT
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.

Hot Area:

Answer Area:

Answer Area

```
{
  "resources": [
    {
      "type": [ "Microsoft.Sql/servers" | "Microsoft.SqlVirtualMachines/sqlVirtualMachines" | "Microsoft.Synapse/workspaces/sqldatabases" ]
      "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')))]"
    …
        }
```

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart

QUESTION 16
HOTSPOT
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.

Hot Area:

Answer Area

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

GO

SELECT SUSER_SNAME()

[          ▼]
REVERT
REVOKE
ROLLBACK

GO
```

Answer Area:

Answer Area

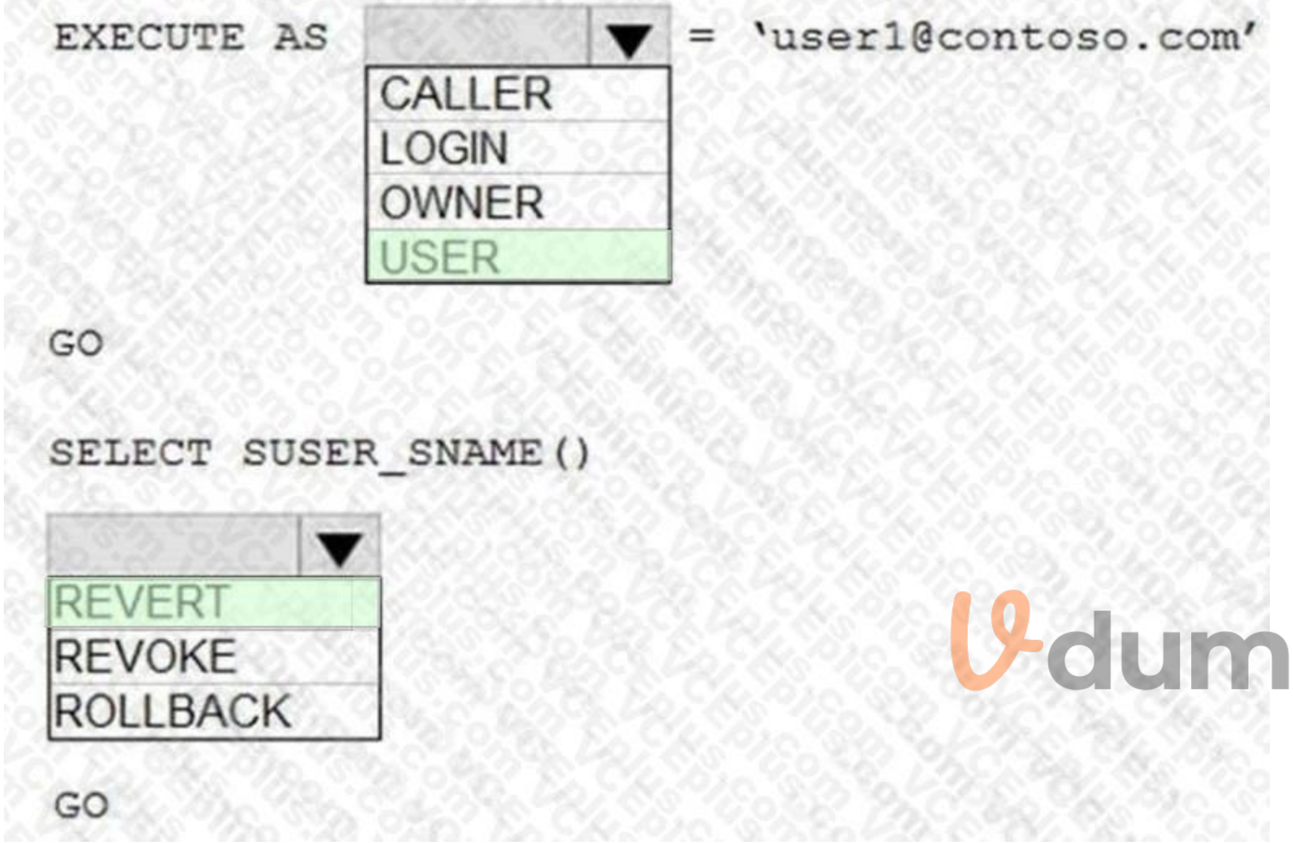

Section:
Explanation:
Reference: https://docs.microsoft.com/en-us/sql/t-sql/statements/execute-as-transact-sql?view=sql-server-ver15 https://docs.microsoft.com/en-us/sql/t-sql/functions/suser-sname-transact-sql?view=sql-server-ver15

QUESTION 17
DRAG DROP
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.

Select and Place:

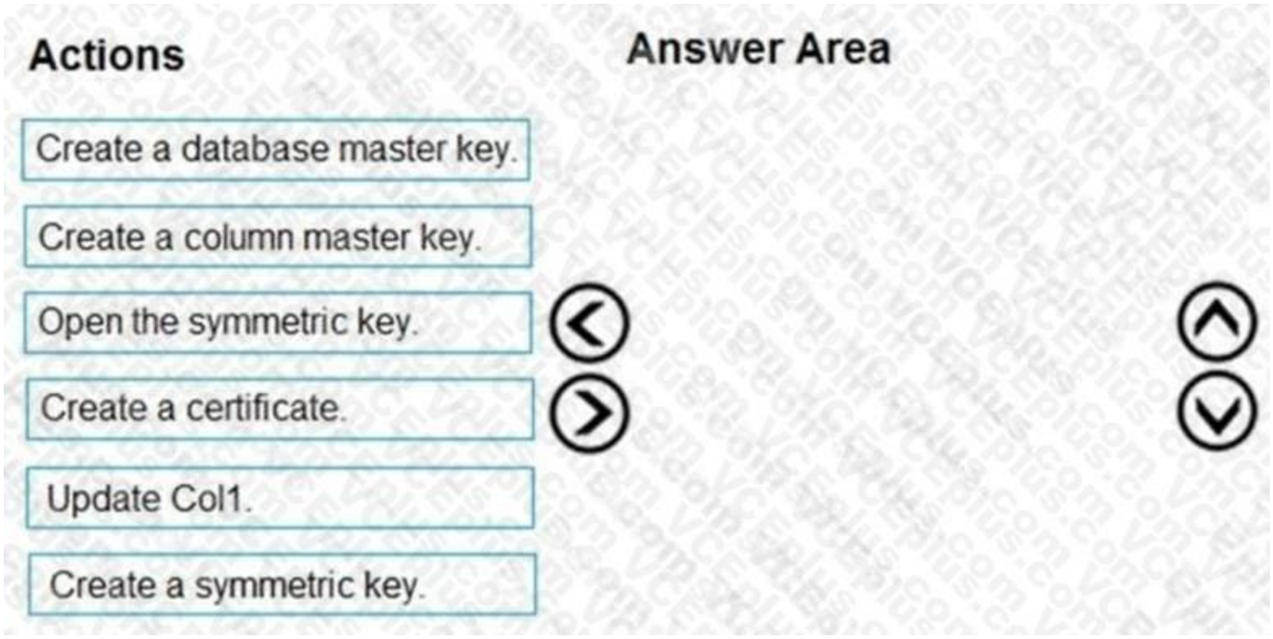

Correct Answer:

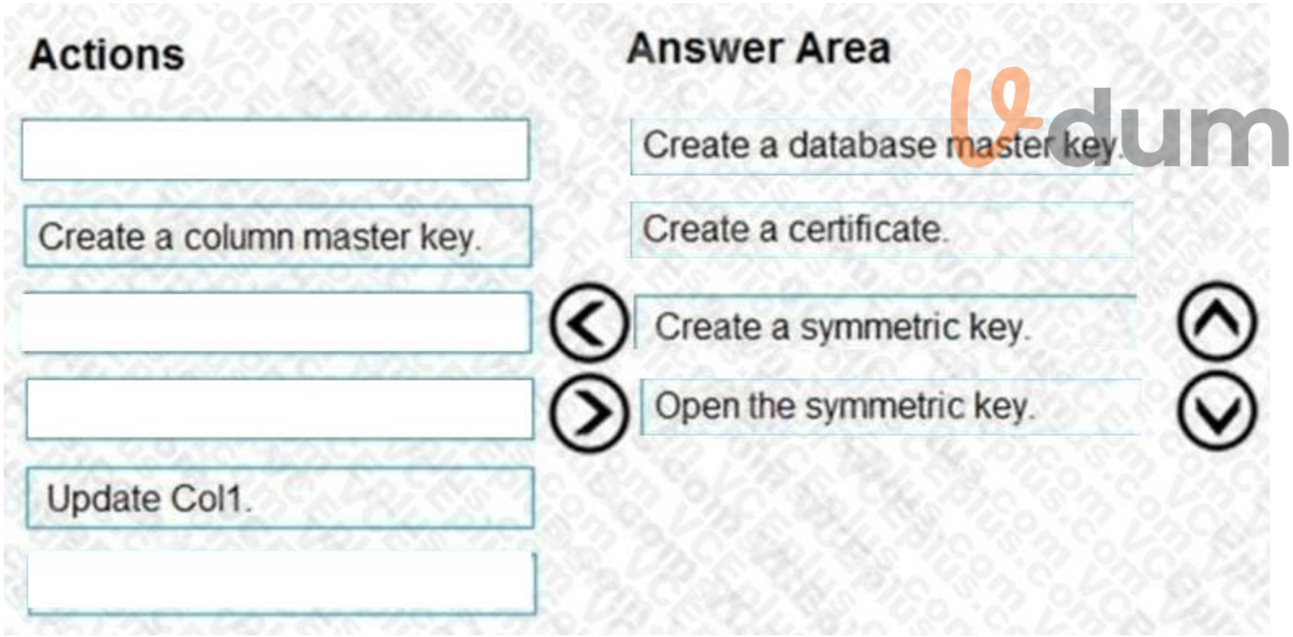

Section:
Explanation:
Reference: https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/

QUESTION 18
HOTSPOT
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.

Hot Area:

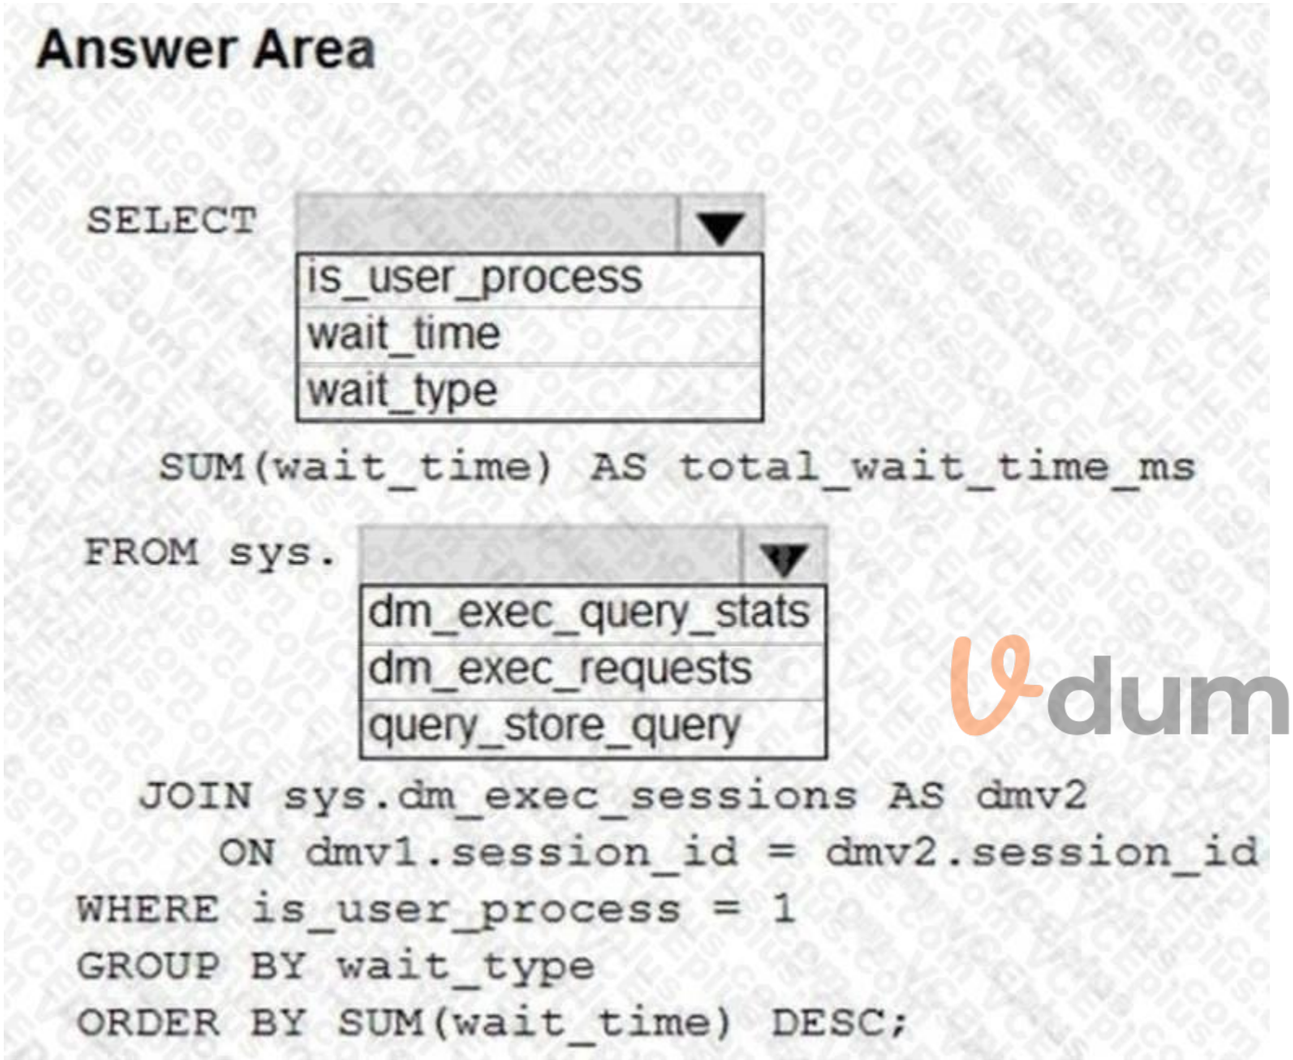

Answer Area:

Answer Area

```
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 dmv2
        ON dmv1.session_id = dmv2.session_id
    WHERE is_user_process = 1
    GROUP BY wait_type
    ORDER BY SUM(wait_time) DESC;
```


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

QUESTION 19
HOTSPOT
You configure a long-term retention policy for an Azure SQL database as shown in the exhibit. (Click the Exhibit tab.)

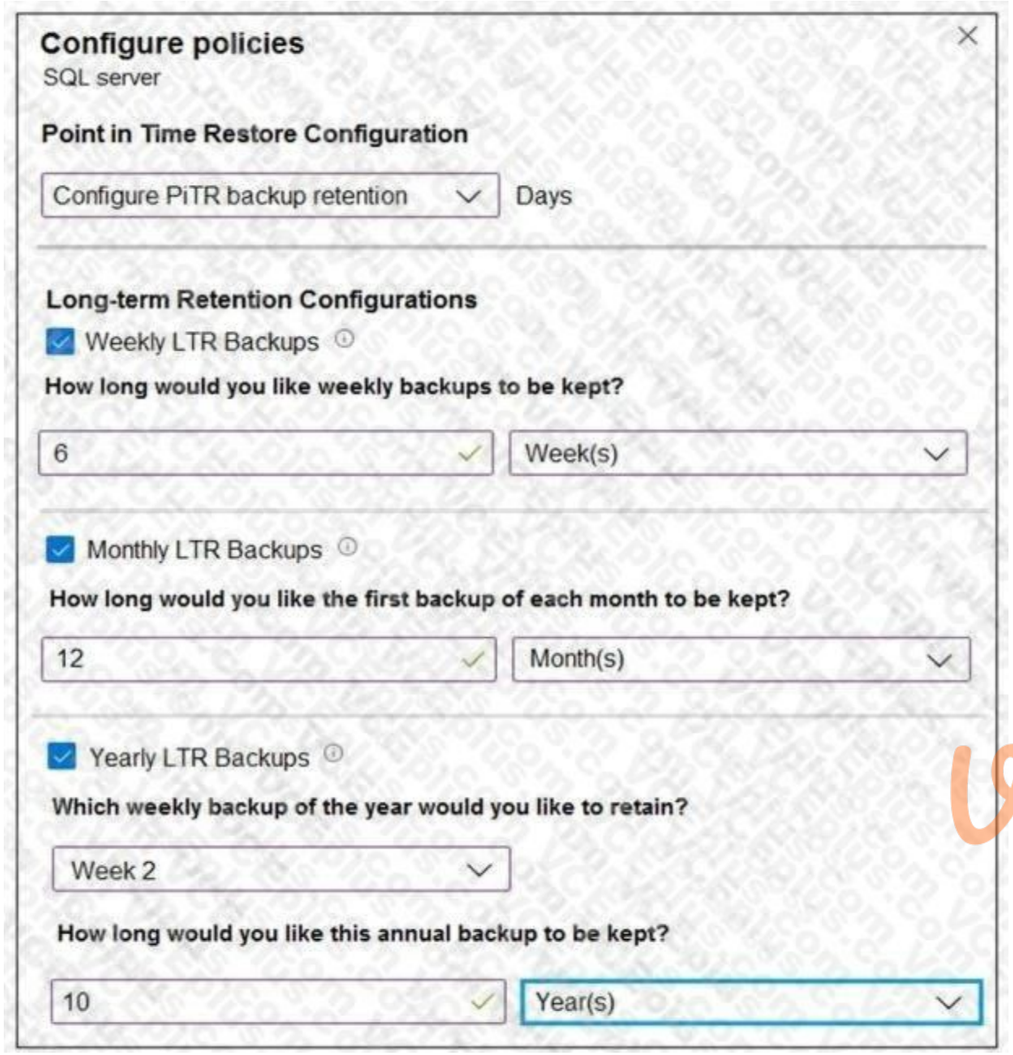

The first weekly backup occurred on January 4, 2020. The dates for the first 10 weekly backups are:
January 4, 2020
January 11, 2020
January 18, 2020
January 25, 2020
February 1, 2020
February 8, 2020
February 15, 2020
February 22, 2020
February 29, 2020 March 7, 2020
Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

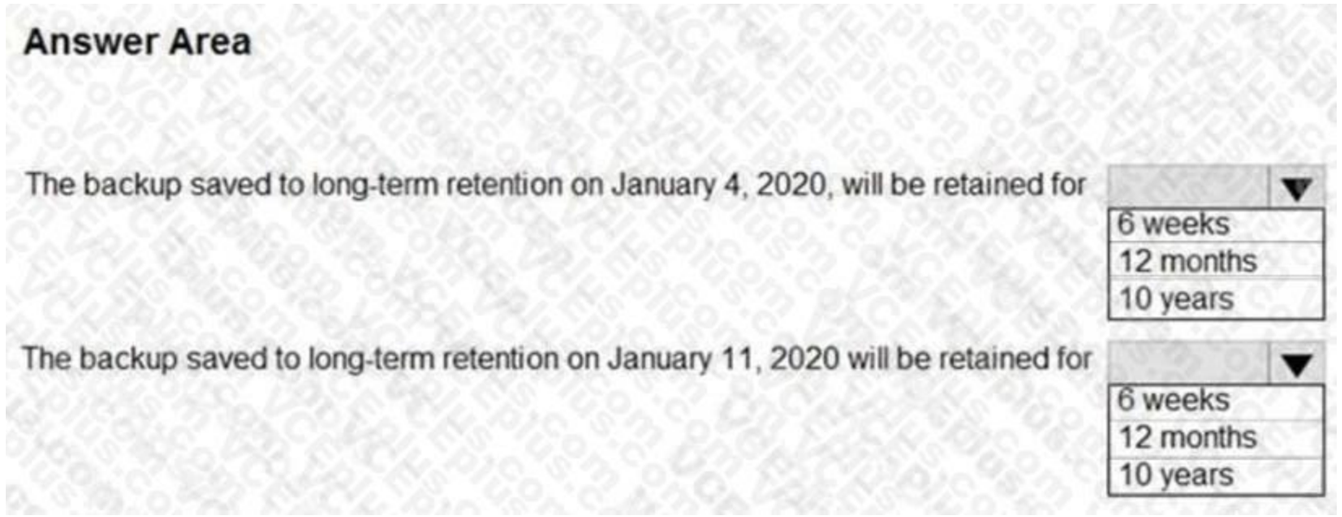

Answer Area:

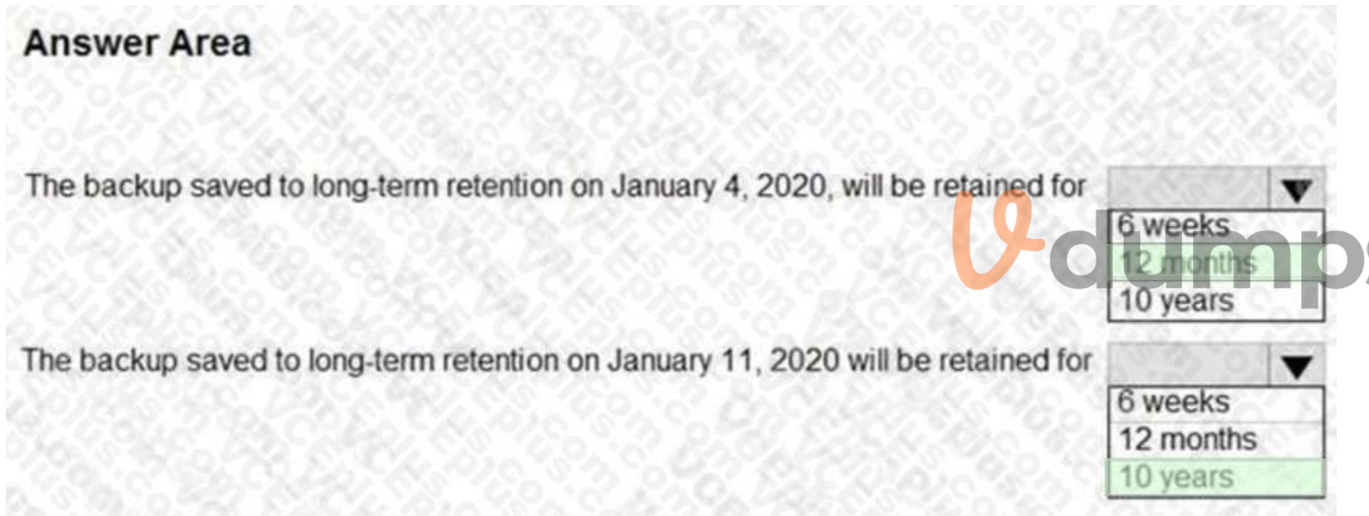

Section:
Explanation:

QUESTION 20
HOTSPOT
You have an Azure SQL database.
You run the following PowerShell script.

```
$serverName = "SERVER1"
$resourceGroup = "RG1"
$dbName = "DB1"

Connect-AzAccount

$server = Get-AzSqlServer -ServerName $serverName -ResourceGroupName
$resourceGroup

Set-AzSqlDatabaseBackupShortTermRetentionPolicy -ResourceGroupName $resourceGroup
-ServerName $server '
    -DatabaseName $dbName -RetentionDays 21

Set-AzSqlDatabaseBackupLongTermRetentionPolicy -ServerName $serverName -
DatabaseName $dbName '
    -ResourceGroupName $resourceGroup -WeeklyRetention P52W -YearlyRetention PSY
-WeekOfYear 52
```

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

Hot Area:

Answer Area

Statements	Yes	No
DB1 can be restored to a specific point in time 30 days ago.	○	○
DB1 can be restored from a weekly backup performed six months ago.	○	○
DB1 can be restored from a yearly backup performed six years ago.	○	○

Answer Area:

Answer Area

Statements	Yes	No
DB1 can be restored to a specific point in time 30 days ago.	○	○
DB1 can be restored from a weekly backup performed six months ago.	○	○
DB1 can be restored from a yearly backup performed six years ago.	○	○

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasebackupshorttermretentionpolicy?view=azps-7.2.0 https://docs.microsoft.com/enus/powershell/module/az.sql/set-azsqldatabasebackuplongtermretentionpolicy?view=azps-7.2.0

QUESTION 21
HOTSPOT
You have an Azure SQL managed instance.
You need to restore a database named DB1 by using Transact-SQL.
Which command should you run? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Section:
Explanation:
https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql?view=azuresqldb-mi-current&preserve-view=true

QUESTION 22
DRAG DROP
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.

Select and Place:

Statements

```
RESTORE DATABASE db2 FROM URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://myssqlbackups.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://myssqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_4.bak'
WITH CREDENTIAL = 'sqlbackup';
GO
```

```
RESTORE DATABASE db2 FROM URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_1.bak'
, URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_2.bak'
,URL =
'https://myssqlbackups.blob.core.windows.net
/backups/db1_3.bak'
, URL =
'https://myssqlbackups.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://myssqlbackups.blob.core.windows.net
/backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET = '<SAS_TOKEN>'
GO
```

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

GO
```

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

Answer Area

Correct 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'
, 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
```

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

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
/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'
```

Section:
Explanation:

QUESTION 23
HOTSPOT
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.

Hot Area:

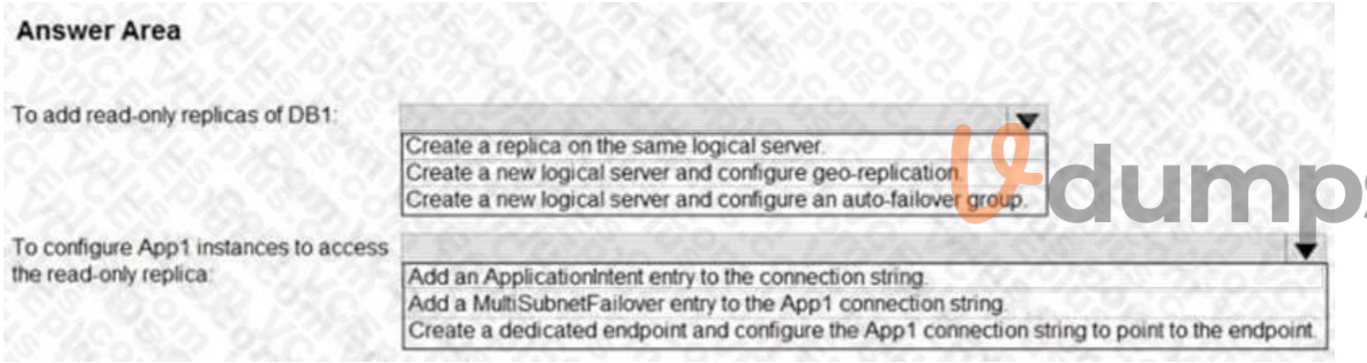

Answer Area:

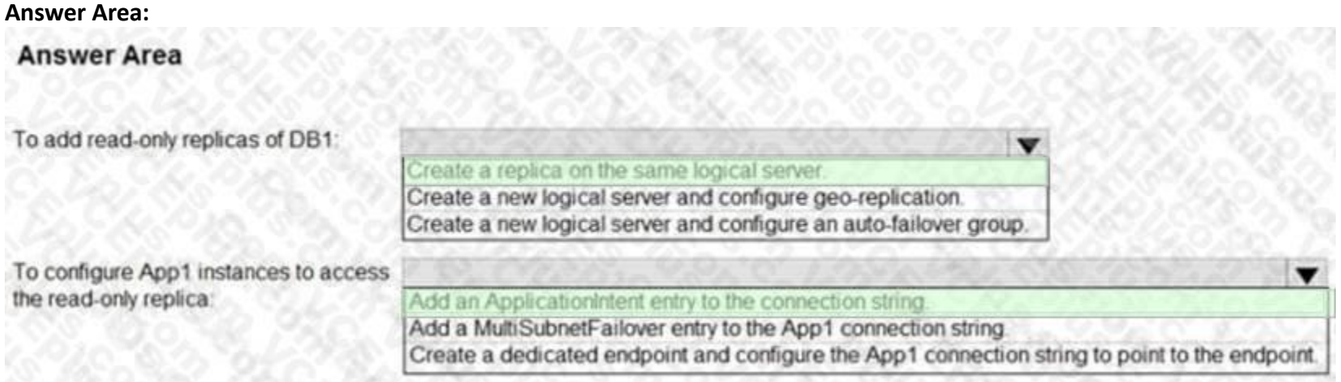

Section:
Explanation:
Reference:
https://sqlserverguides.com/read-only-replica-azure-sql/

QUESTION 24

HOTSPOT
You have a SQL Server on Azure Virtual Machines instance named VM1 that hosts a database named DB1. You run the following query.

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

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

Hot Area:

Answer Area

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

Answer Area:

Answer Area

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

Section:
Explanation:

QUESTION 25
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

Correct Answer: C
Section:

QUESTION 26
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

You need to configure a connection between VM1 and MIL The solution must meet the following requirements:
• The connection must be encrypted.
• Network latency must be minimized.
What should you implement?

A. virtual network peering
B. private endpoints
C. service endpoints
D. a site-to-site VPN

Correct Answer: B
Section:

QUESTION 27
DRAG DROP
You have an Azure subscription that contains an Azure SQL managed instance, a database named db1, and an Azure web app named Appl. Appl uses db1. You need to enable Resource Governor for a App1. The solution must meet the following requirements:
App1 must be able to consume all available CPU resources.
App1 must have at least half of the available CPU resources always available.
Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Correct Answer:

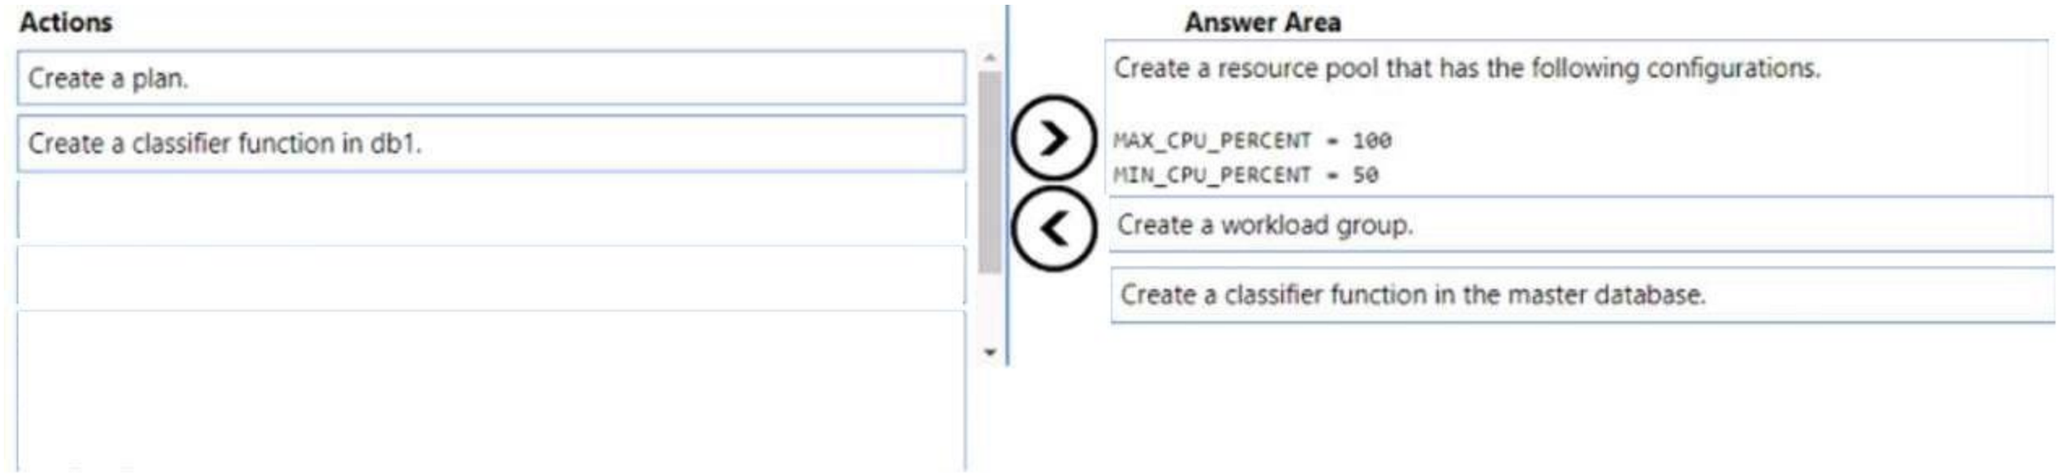

Section:
Explanation:

QUESTION 28
HOTSPOT
You have an Azure SQL database named that contains a table named Table1.
You run a query to bad data into Table1.
The performance Of Table1 during the load operation are shown in exhibit.

Hot Area:

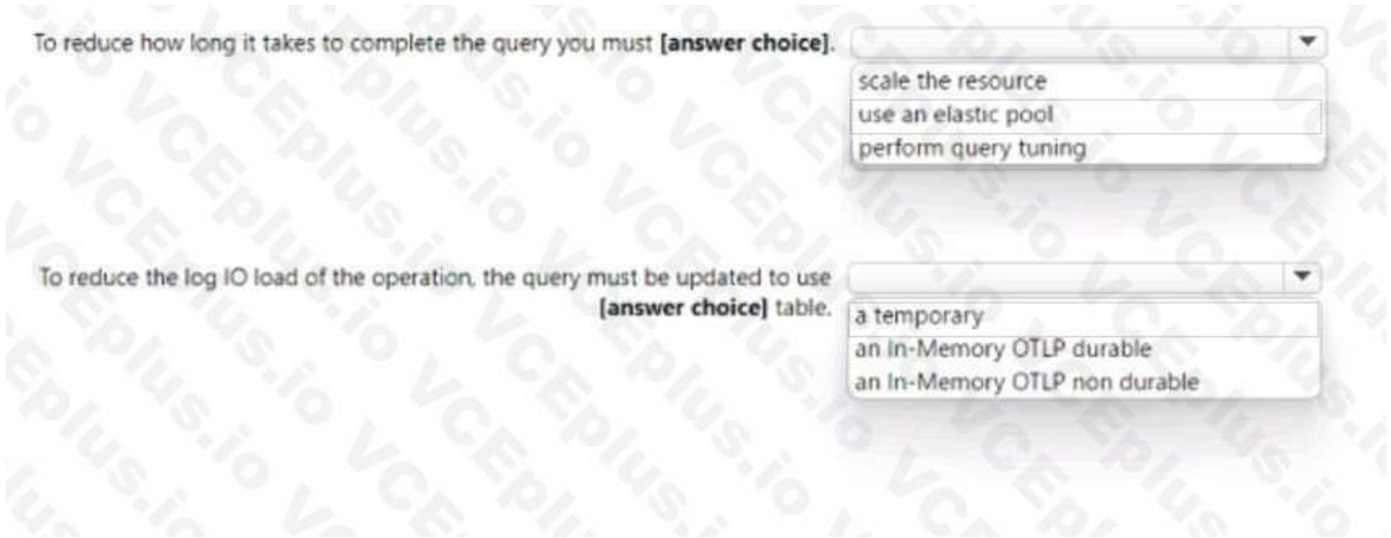

Answer Area:

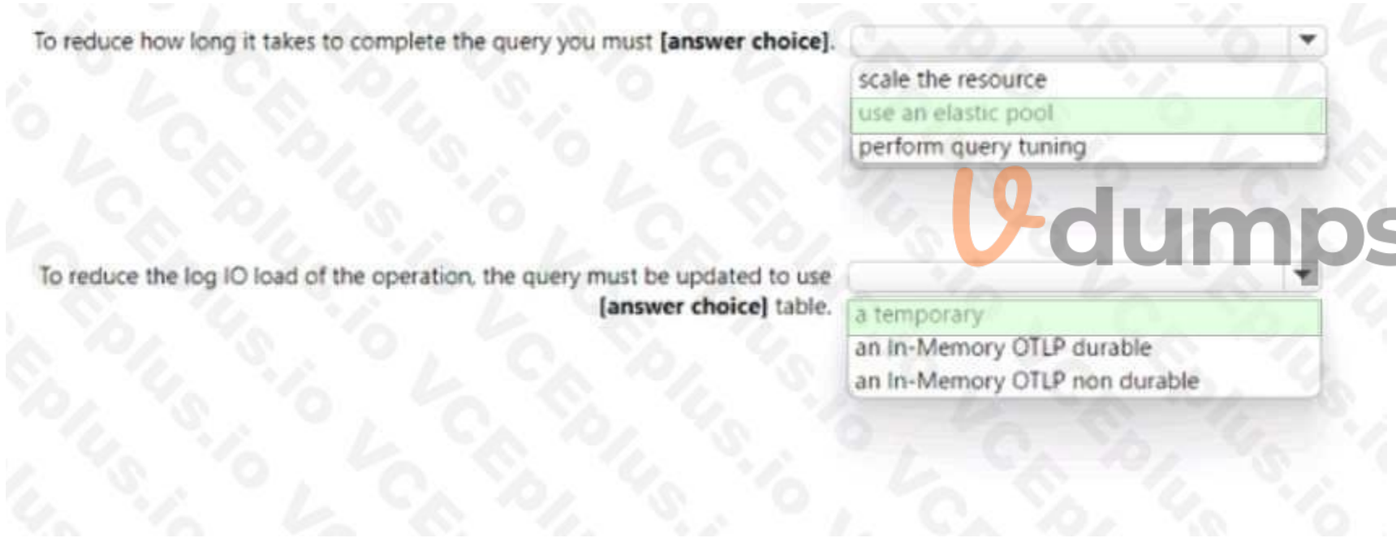

Section:
Explanation:

QUESTION 29
DRAG DROP
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)

Source spid1595

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

You need to ensure That db1 can process transactions.

Select and Place:

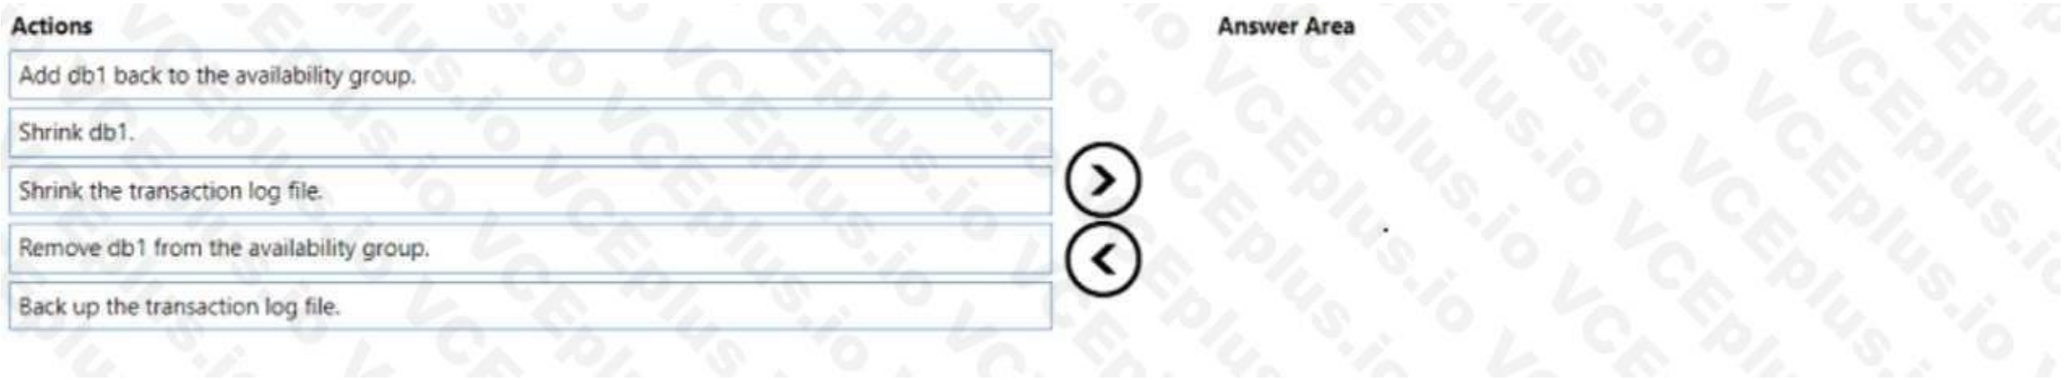

Correct Answer:

Section:
Explanation:

QUESTION 30
HOTSPOT
You have an Azure SQL database.
You need to identify whether a delayed query execution is associated to a RESOURCE wait.
How should you complete the Transact –SQL statement? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
SELECT [ wait_type | context_info | wait_resource ]
    SUM(wait_time) AS total_wait_time_ms
FROM [ sys.dm_exec_requests | sys.dm_exec_connections | sys.dm_db_partition_stats ] AS dmv1
 JOIN svs.dm exec sessions AS sess
     ON dmv1.session id = sess.session id
    WHERE is_user_process = 1
    GROUP BY TARGET1
    ORDER BY SUM(wait_time) DESC;
```

Answer Area:

Answer Area

```
SELECT [wait_type | context_info | wait_resource]
    SUM(wait_time) AS total_wait_time_ms
FROM [sys.dm_exec_requests | sys.dm_exec_connections | sys.dm_db_partition_stats] AS dmv1
  JOIN svs.dm exec sessions AS sess
      ON dmv1.session id = sess.session id
    WHERE is_user_process = 1
    GROUP BY TARGET1
    ORDER BY SUM(wait_time) DESC;
```

Section:
Explanation:

QUESTION 31
You have an Azure subscription that contains two Azure SQL managed instances named SQLMI1 and SQLMI2 . SQLM2 contains a database named DB1 and a user named User1. User1 drops DB1.
You need to perform a point-in-time restore of DB1 to SQLMI2.

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

Correct Answer: C
Section:

QUESTION 32
You have an Azure subscription that contain 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 Workspace1.

A. Create the private endpoint connection on SQLMI1.

B. Configure Azure SQL Analytics to use Workspace1.
C. Modify the Computer + storage settings for SQLMI1.
D. Modify the diagnostic settings for SQLMI1.

Correct Answer: D
Section:

QUESTION 33
DRAG DROP
You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored procedures. You need to ensure that SOL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

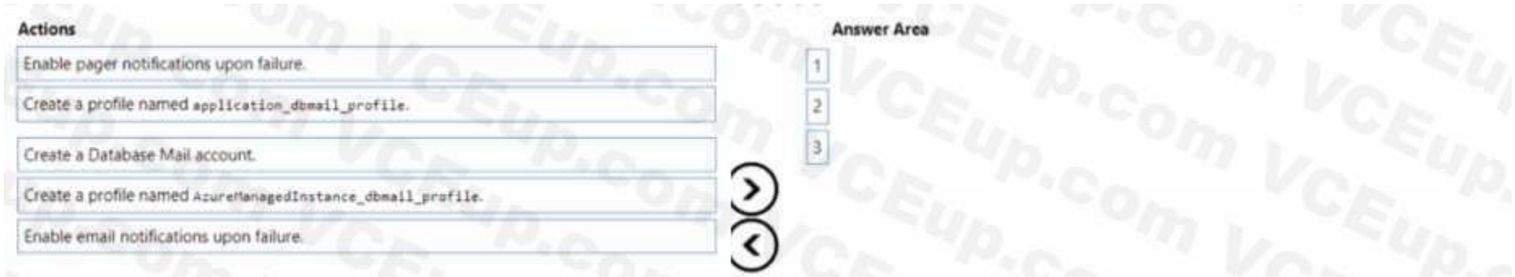

Correct Answer:

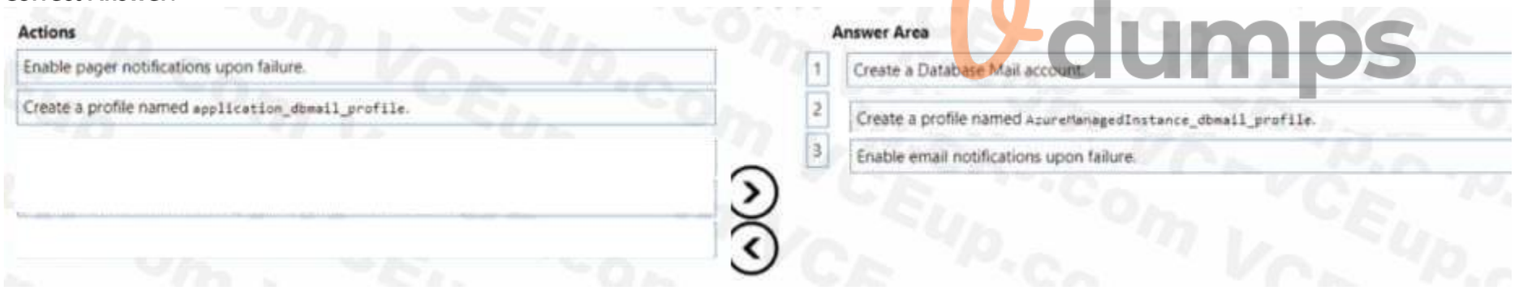

Section:
Explanation:

QUESTION 34
You have An Azure SQL managed instance.
You need to configure the SQL Server Agent service to email job notifications.
Which statement should you execute?

A.
```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'sysadmin_dbmail_profile';
```
B.
```
EXECUTE msdb.dbo.sysmail_edd_profile_sp @profile_name = 'application_dbmail_profile';
```
C.
```
EXECUTE msdb.dbo.sysmail_add_profile_sp @profile_name = 'AzureManagedInstance_dbmail_profile';
```
D.

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

Correct Answer: B
Section:

QUESTION 35
HOTSPOT
You have an Azure SQL database named D61.
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 m the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

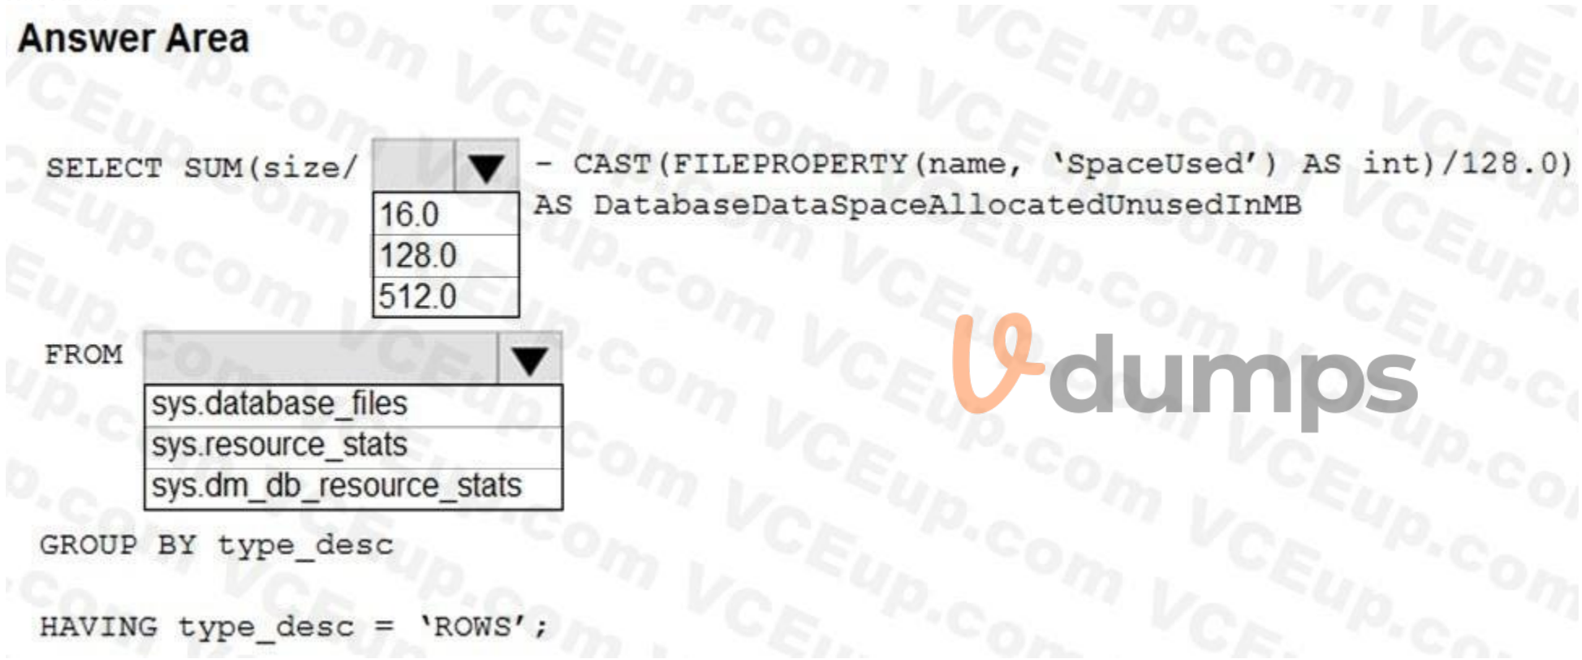

Answer Area:

Answer Area

```
SELECT SUM(size/ [ ▼ ] - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
                                AS DatabaseDataSpaceAllocatedUnusedInMB
FROM [ ▼ ]
GROUP BY type_desc

HAVING type_desc = 'ROWS';
```

16.0
128.0
512.0

sys.database_files
sys.resource_stats
sys.dm_db_resource_stats

Section:
Explanation:

QUESTION 36
HOTSPOT
You configure backup for an Azure SQL database as shown in the following exhibit.

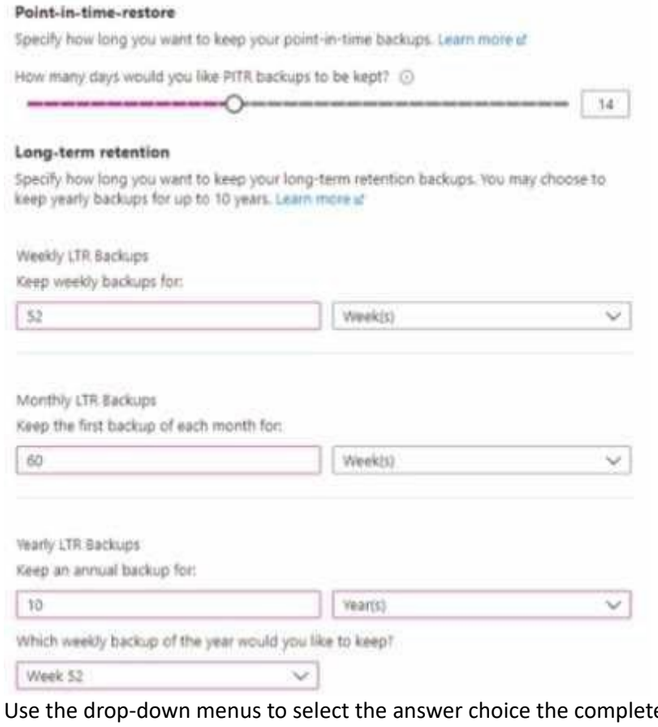

Use the drop-down menus to select the answer choice the completes each statement based on the

information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

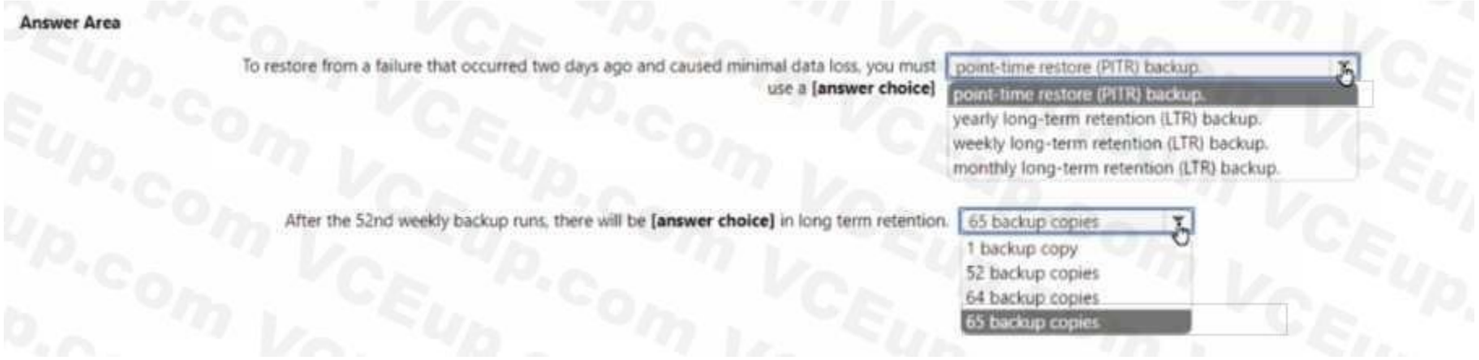

Answer Area:

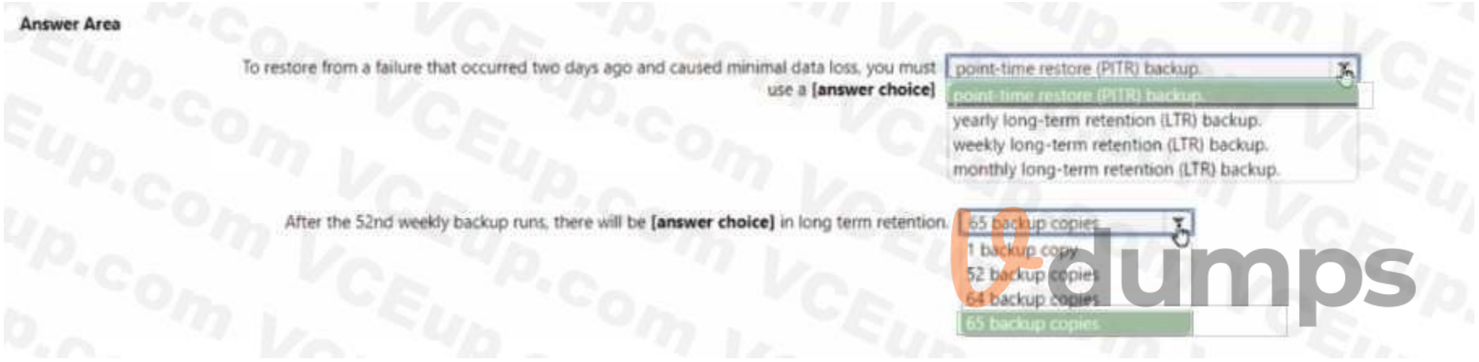

Section:
Explanation:

QUESTION 37
HOTSPOT
You have an Azure subscription that contains a group named Group1 and an Azure SQL managed instance that hosts a database named 081. You need to ensure that Group 1 has read access to new tables created m 06I The solution must use the principle of least privilege How should you complete the Transact-SQL statement' To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 38
You have an Azure SQL database named DB1.
A user named User 1 has an Azure AD account.
You need to provide User1 with the ability to add and remove columns from the tables m DBV The solution must use the principle of least privilege.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point

A. Assign the database user the db.ddladmm role.
B. Assign the database user the db.owner role.
C. Create a contained database user.
D. Create a login and an associated database user.

Correct Answer: A, D
Section:

QUESTION 39
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 need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:
* Use block blob storage.
* Maximize security.

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

Correct Answer: B
Section:

QUESTION 40
You have an Azure AD tenant and a logical Microsoft SQL server named SQL1 that hosts several Azure SQL databases. You plan to assign Azure AD users permissions to the databases automatically by using Azure Automation.
You need to create the required Automation accounts.
Which two accounts should you create? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. From the Azure Active Directory admin center, create a service principal.
B. From the Azure Active Directory admin center, create a user-assigned managed identity for SQL1.

C. On SQL1. create a SQL user in the databases.
D. On SQL1, create a SQL login.
E. From the Azure Active Directory admin center, create an external identity.

Correct Answer: A, C
Section:

QUESTION 41
You have an Azure SQL database named DB1.
You need to encrypt DB1. The solution must meet the following requirements;
- Encrypt data in motion.
- Support comparison operators.
- Provide randomized encryption.
What should you include in the solution?

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

Correct Answer: A
Section:

QUESTION 42
You have a SQL Server on Azure Virtual Machines instance named SQLVM1 that was deployed by using an Azure Marketplace SQL Server 2019 Enterprise image. You need to change the Microsoft SQL Server instance on SQLVM1 to the Standard edition. The solution must ensure licensing compliance. What should you do first?

A. From the SQL Server Installation Center on SQLVM1. run the Edition Upgrade wizard.
B. From SQLVM1, uninstall the SQL Server instance.
C. From the SQL Server Installation Center on SQLVM1. run the Repair wizard.
D. From the Azure portal, reconfigure SQLVM1.

Correct Answer: B
Section:

QUESTION 43
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. database mirroring
B. distributed availability groups
C. Always On Availability Group
D. Azure Database Migration Service

Correct Answer: D
Section:

QUESTION 44
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

Correct Answer: B
Section:

QUESTION 45
You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database. The solution must meet the following requirements:
- Support real-time data replication to a different geographic region.
- Use Azure as a disaster recovery target.
- Minimize costs and administrative effort.
What should you include in the recommendation?

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

Correct Answer: D
Section:

QUESTION 46
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

Correct Answer: D
Section:

QUESTION 47
You have an instance of SQL Server on Azure Virtual Machines.
You need to ensure that a user named User1 can configure proxy accounts for SQL Server Agent jobs.
The solution must use the principle of least privilege.
Which role should you assign to User1?

A. sysadmin
B. SQLAgentUserRole
C. SQLAgentReaderRole
D. SQLAgentOperatorRole

Correct Answer: A
Section:

QUESTION 48
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

Correct Answer: A, D
Section:

QUESTION 49
You have a single availability set that contains two SQL Server on Azure Virtual Machines instances. The instances were deployed by using an Azure Marketplace SQL Server 2019 Enterprise image that has the latest cumulative updates applied. The instances are configured as the nodes of a failover cluster instance (FCI) named FCI1. You need to ensure that client applications can connect to FCI1. The solution must meet the following requirements:
- Provide an availability SLA
- Minimize costs.
What should you create?

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

Correct Answer: C
Section:

QUESTION 50
You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group. You need to recommend a deployment solution that meets the

following requirements:
- Provides a Service Level Agreement (SLA) of at least 99.95%
- Replicates databases in the same group synchronously
- Minimizes the latency of database writes
What should you recommend?

A. Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group.
B. Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group.
C. Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group.
D. Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group.

Correct Answer: A
Section:

QUESTION 51
HOTSPOT
You have an on-premises Microsoft SQL Server 2016 instance that hosts a database named db1. You have an Azure subscription that contains an Azure SQL managed instance named Mil. You plan to perform an online migration of db1 to MM by using Azure Database Migration Service.
You need to create the backups for the migration. The solution must minimize the number of backup files created. Which type of backups should you create, and how should you store the backups? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

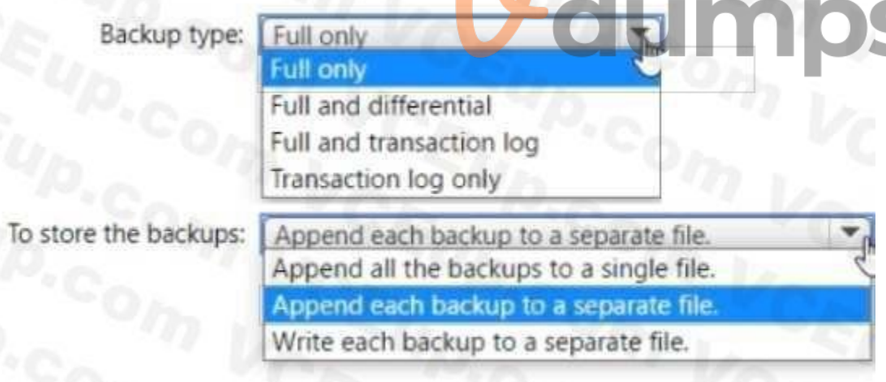

Answer Area:

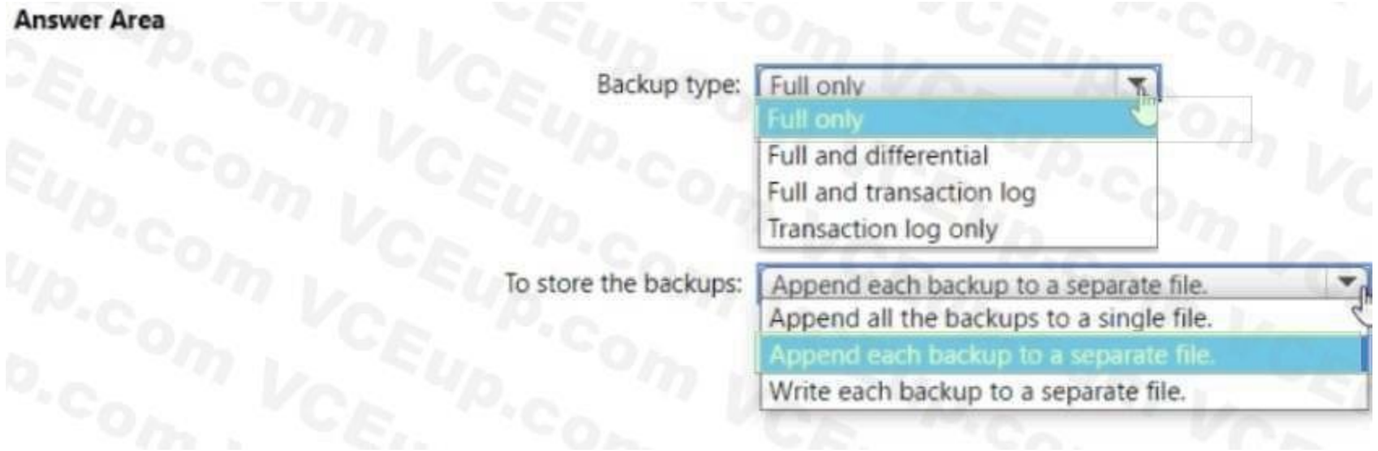

Section:
Explanation:

QUESTION 52
HOTSPOT
You have an Azure subscription that contains a resource group named RG1. RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL You need to use PowerShell to enable and configure automated patching for SQL The solution must include both SQL Server and Windows security updates.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Answer Area

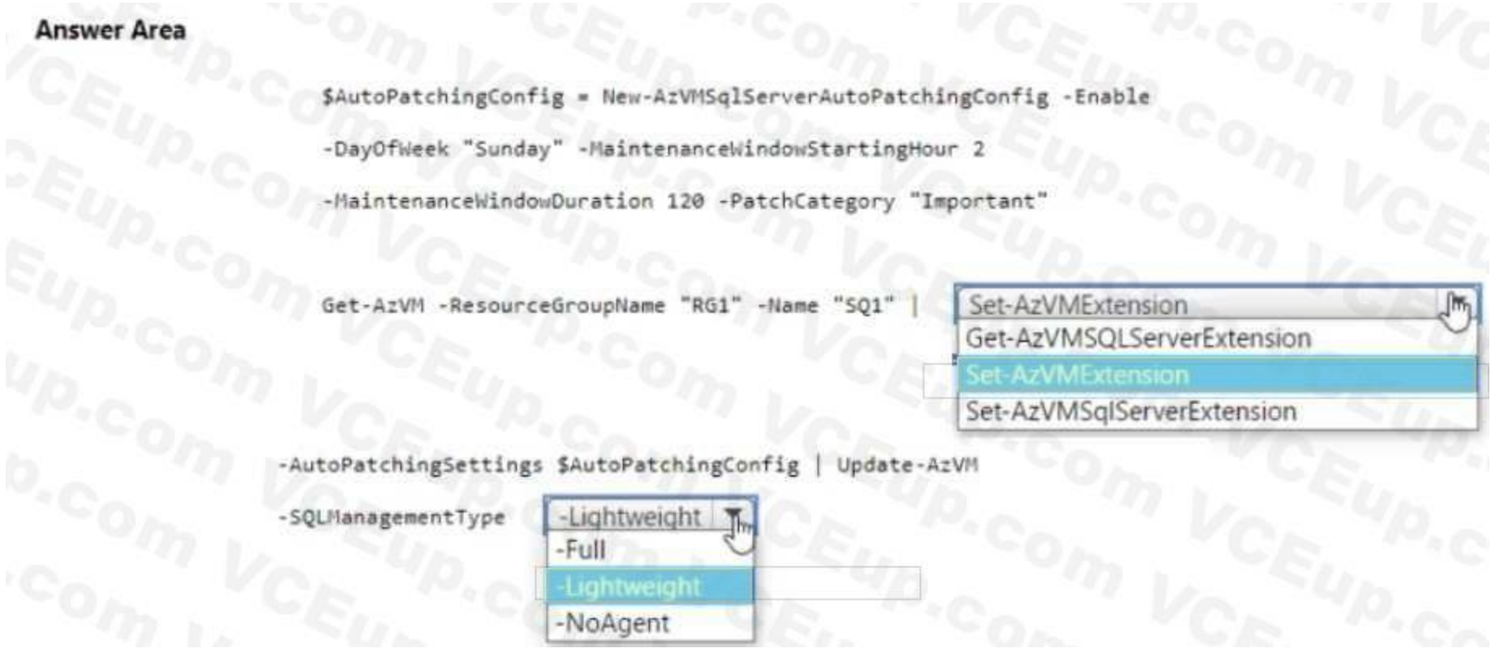

Section:
Explanation:

QUESTION 53
You have an Azure SQL database.
You need to implement a disaster recovery solution that meets the following requirements:
• Minimizes how long it takes to recover the database if a datacenter fails
• Minimizes administrative effort
What should you include in the solution?

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

Correct Answer: C
Section:

QUESTION 54
HOTSPOT
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.

Hot Area:

Answer Area

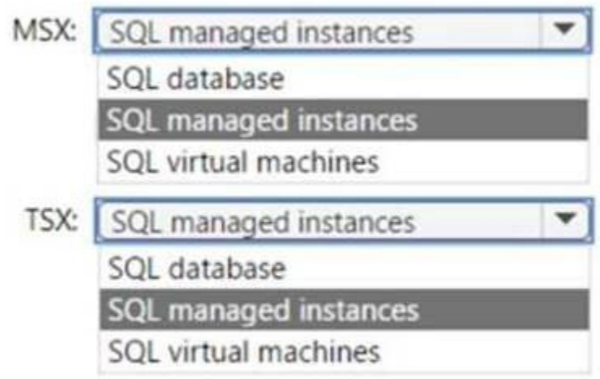

Answer Area:

Answer Area

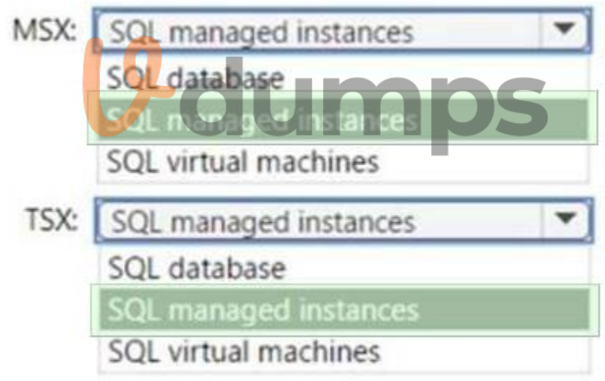

Section:
Explanation:

QUESTION 55
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 User! can create a scheduled task to perform a full backup of DB1. The solution must use the principle of least privilege.
Which built-in database role should you assign to User1?

A. SQLAgentReaderRole
B. db.owner

C. SQLAgentOperatorRole
D. SQLAgentUserRole

Correct Answer: C
Section:

QUESTION 56
You have an on-premises datacenter that contains a 2-TB Microsoft SQL Server 2019 database named DB1.
You need to recommend a solution to migrate DB1 to an Azure SQL managed instance. The solution must minimize downtime and administrative effort.
What should you include in the recommendation?

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

Correct Answer: B
Section:

QUESTION 57
You have an Azure subscription that contains 50 instances of SQL Server on Azure Virtual Machines. The instances host 500 Azure SQL databases. You need to ensure that all the databases have the same configuration. The solution must meet the following requirements:
• Auditing must be enabled.
• Azure Defender must be enabled.
• Public network access must be disabled.
• Administrative effort must be minimized.
Which two resources should you create in the subscription? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Correct Answer: A, B
Section:

QUESTION 58
HOTSPOT
You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD sen/ice principal named appl.
You need to ensure that appl can access db1. The solution must use the principle of least privilege.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

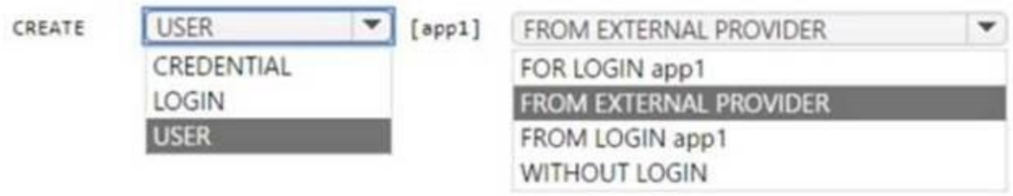

Answer Area:

Answer Area

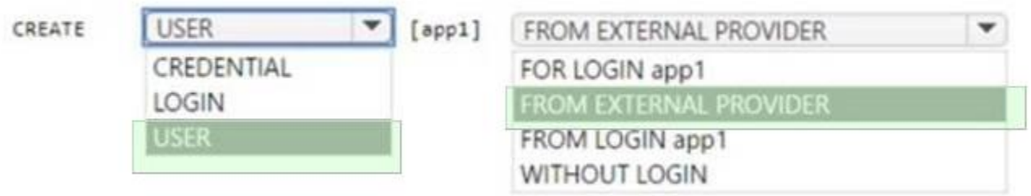

Section:
Explanation:

QUESTION 59
HOTSPOT
You have an Azure subscription.
You need to deploy an Azure SQL managed instance that meets the following requirements:
• Optimize latency.
• Maximize the memory-to-vCore ratio.
Which service tier and hardware generation should you use? To answer, select the apocopate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

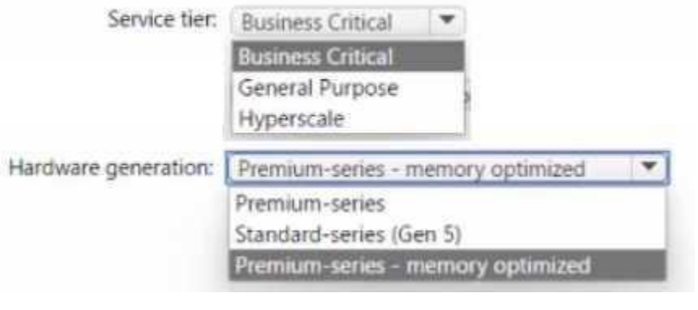

Answer Area:

Answer Area

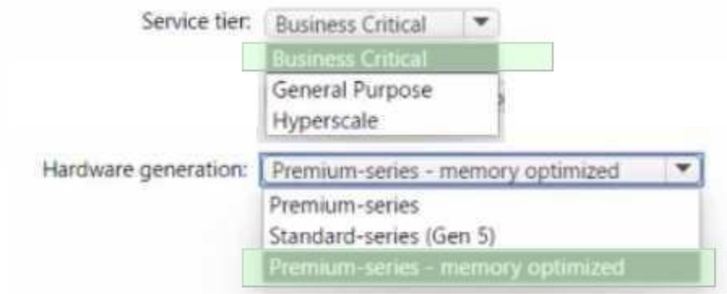

Section:
Explanation:

QUESTION 60
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 SQlDC1.
Which type of replication should you use?

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

Correct Answer: C
Section:

QUESTION 61
HOTSPOT
You have an Azure subscription.
You need to deploy a logical SQL server by using an Azure Resource Manager (ARM) template. The solution must ensure that the server will allow inbound connectivity from any Azure resource.
How should you complete the template? To answer, select the appropriate options in the answer area
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 62
You have an Azure subscription that contains the resources shown in the following table.

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

You plan louse SQLDB11 as an elastic job database to run jobs on SQlDb11 andSOtDB22. What is the minimum number of database scoped credentials required tor the elastic jobs?

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

Correct Answer: A
Section:

QUESTION 63
You have an Azure subscription.
You need to deploy an Instance of SQL Server on Azure Virtual Machines. The solution must meet the following 'requirements:
* Custom performance configuration. such as lCPS. capacity, and throughout, must be supported.
* Costs must be minimized
Which type of disk should you include in the solution?

A. Premium SSD v2
B. Premium SSD
C. Ultra SSD
D. Standard SSD

Correct Answer: A
Section:

QUESTION 64
You have an Azure subscription.
You create a logical SQL server that hosts four databases Each database will be used by a separate customer.
You need to ensure that each customer can access only its own database. The solution must minimize administrative effort
Which two actions should you perform? Each correct answer presents part of the solution
NOTE: Each correct selection is worth one point.

A. Create a network security group (NSG)
B. Create a server-level firewall rule
C. Create a private endpoint
D. Create a database-level firewall rule.
E. Deny public access.

Correct Answer: C, D
Section:

QUESTION 65
DRAG DROP
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.

Select and Place:

Objects

- a database encryption key in the master database
- a master key in DB1
- a certificate in DB1
- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

Answer Area

Correct Answer:

Objects

- a database encryption key in the master database
- a master key in DB1
- a certificate in DB1

Answer Area

- a master key in the master database
- a certificate in the master database
- a database encryption key in DB1

Section:
Explanation:

QUESTION 66
You have an Azure SQL database named DB1 in the General Purpose service tier.
The performance metrics for DB1 are shown in the following exhibit.

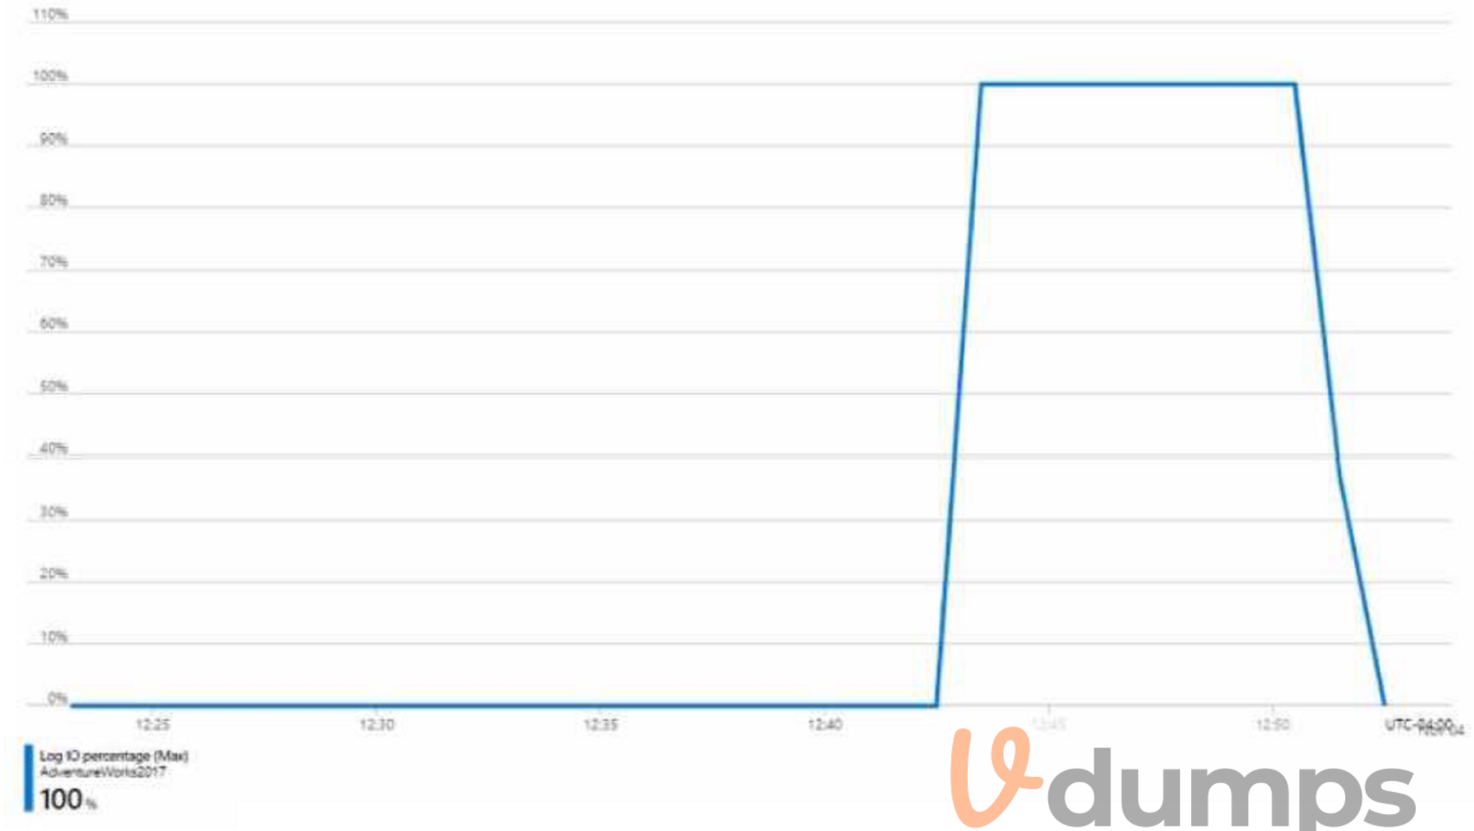

You need to reduce the Log 10 percentage. The solution must minimize costs.
What should you do?

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

Correct Answer: D
Section:

QUESTION 67
You have an Azure SQL Database elastic pool that contains 10 databases.
You receive the following alert.
Msg 1132, Level 16, State 1, Line 1
The elastic pool has reached its storage limit. The storage used for the elastic pool cannot exceed (76800) MBs.
You need to resolve the alert. The solution must minimize administrative effort.
Which three actions can you perform? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Delete data from a database.
B. Remove a database from the pool.

C. Increase the maximum storage of the elastic pool.
D. Shrink individual databases.
E. Enable data compression.

Correct Answer: B, C, D
Section:

QUESTION 68
HOTSPOT
You configure backups for an Azure SQL database as shown in the following exhibit.

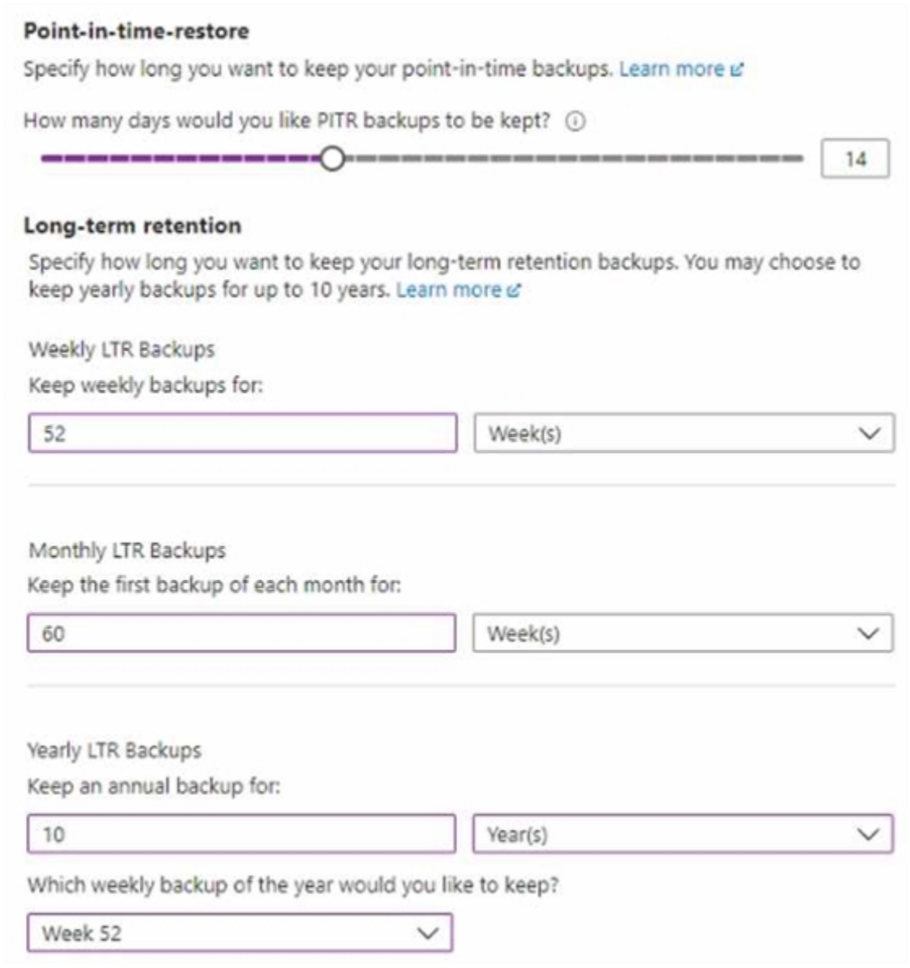

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

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

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

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

Answer Area:

Answer Area

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

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

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

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

Section:
Explanation:

QUESTION 69
You have an Azure subscription that contains two instances of SQL Server on Azure Virtual Machines named VM1 and VM2. Both instances run Microsoft SQL Server 2019 CU8. You need to deploy a failover cluster instance (FCI) to VM1 and VM2. The solution must eliminate the need for the following:
• A distributed network name (DNN)
• A load balancer
What should you do?

A. Deploy VM1 and VM2 to a single proximity placement group.
B. Deploy VM1 and VM2 to different proximity placement groups in the same Azure region.
C. Connect VM1 and VM2 to a single subnet.
D. Connect VM1 and VM2 to different subnets on a single virtual network.

Correct Answer: D
Section:

QUESTION 70
HOTSPOT
You have an Azure subscription that contains a storage account named databasebackups.
You have an Azure SQL managed instance named DB1.
You need to back up DB1 to databasebackups.
How should you complete the commands? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

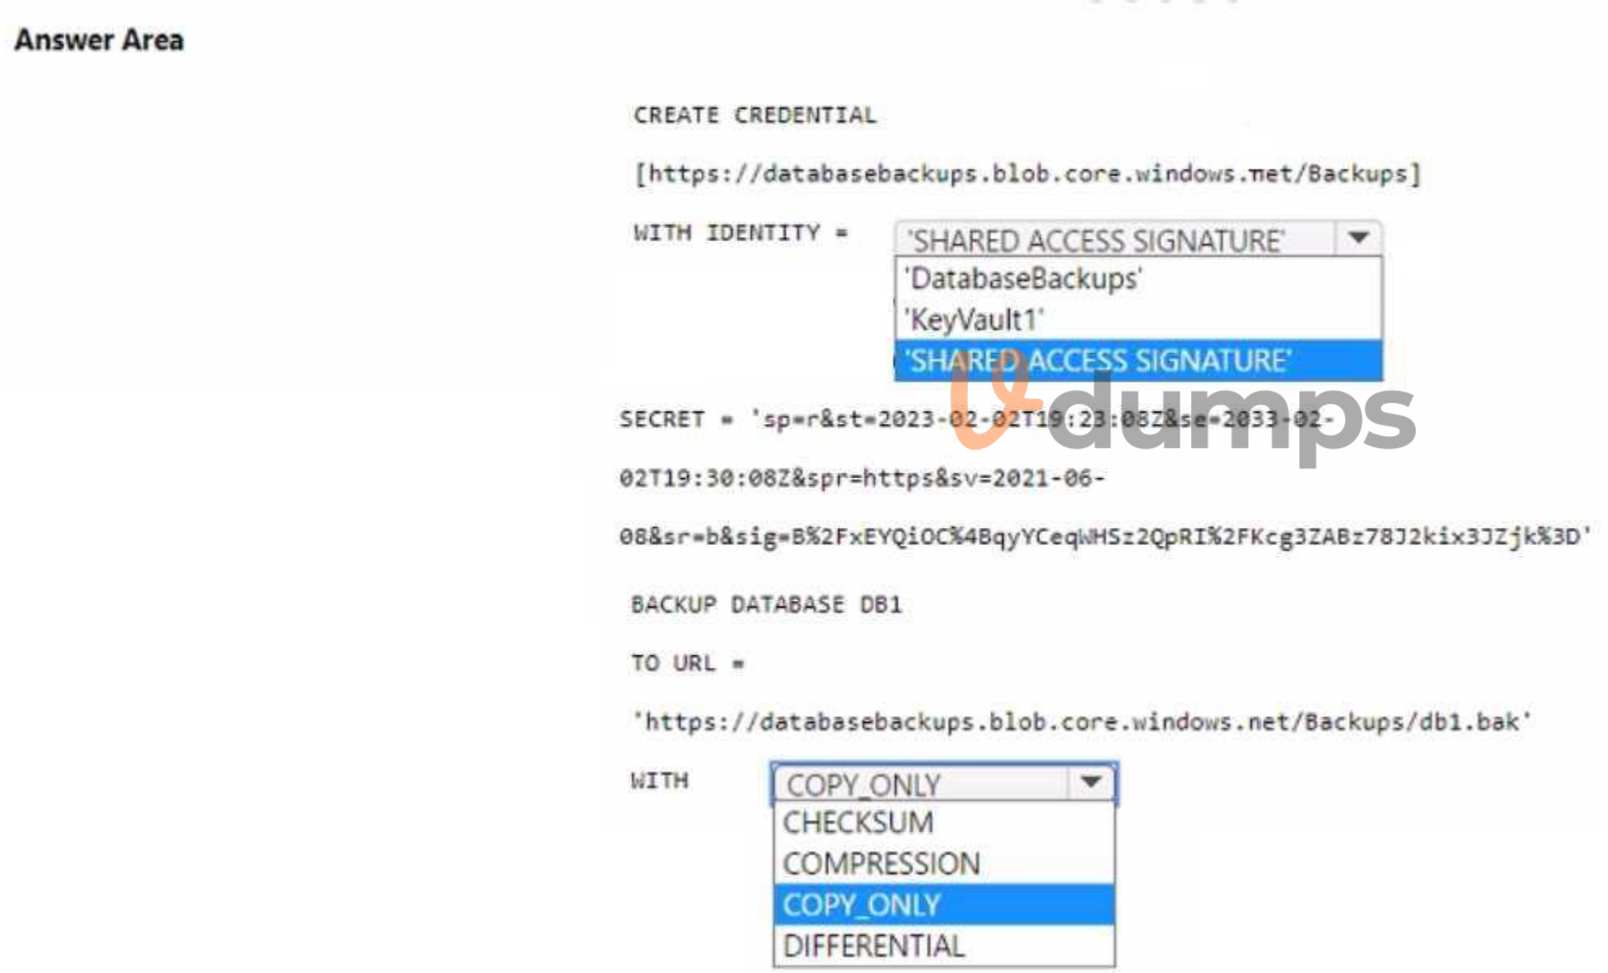

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 71

HOTSPOT

From a website analytics system, you receive data extracts about user interactions such as downloads, link clicks, form submissions, and video plays.

The data contains the following columns:

You need to design a star schema to support analytical queries of the data. The star schema will contain four tables including a date dimension.

To which table should you add each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

EventCategory:
- DimChannel
- DimDate
- DimEvent
- FactEvents

ChannelGrouping:
- DimChannel
- DimDate
- DimEvent
- FactEvents

TotalEvents:
- DimChannel
- DimDate
- DimEvent
- FactEvents

Answer Area:

EventCategory:
- DimChannel
- DimDate
- DimEvent
- FactEvents

ChannelGrouping:
- DimChannel
- DimDate
- DimEvent
- FactEvents

TotalEvents:
- DimChannel
- DimDate
- DimEvent
- FactEvents

Section:
Explanation:

QUESTION 72
HOTSPOT
You have an Azure subscription that contains a logical SQL server. The server hosts two databases named db1 and db2 and an Azure AD service principal named app1.
You need to ensure that app1 can access db1. The solution must use the principle of least privilege.
How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

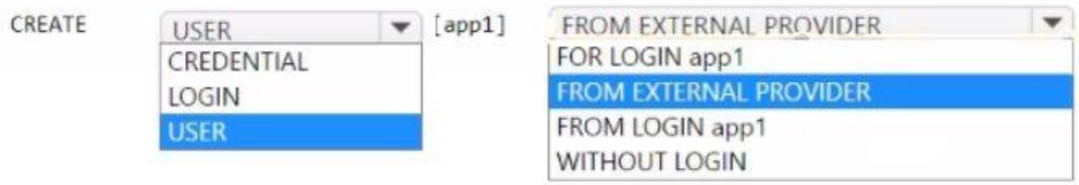

Answer Area:

Answer Area

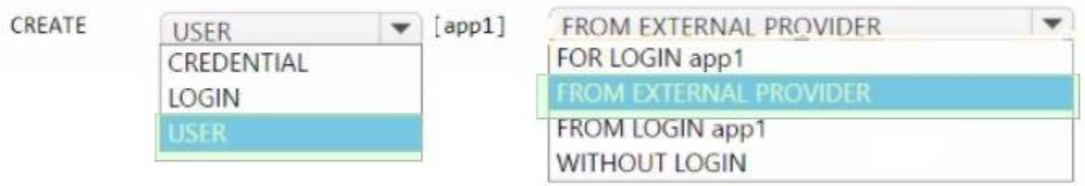

Section:
Explanation:

Answer Area

https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql

QUESTION 73
HOTSPOT
You have an instance of SQL Server on Azure Virtual Machines named VM1.
You need to use an Azure Automation runbook to initiate a SQL Server database backup on VM1.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

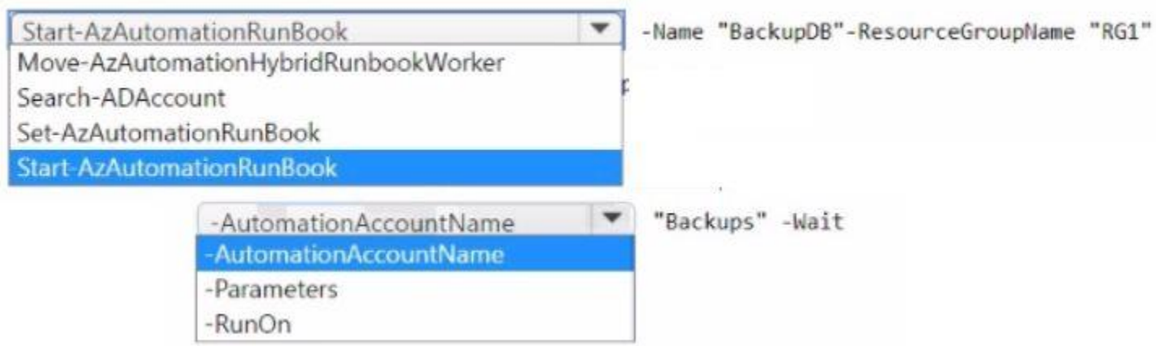

Answer Area:

Answer Area

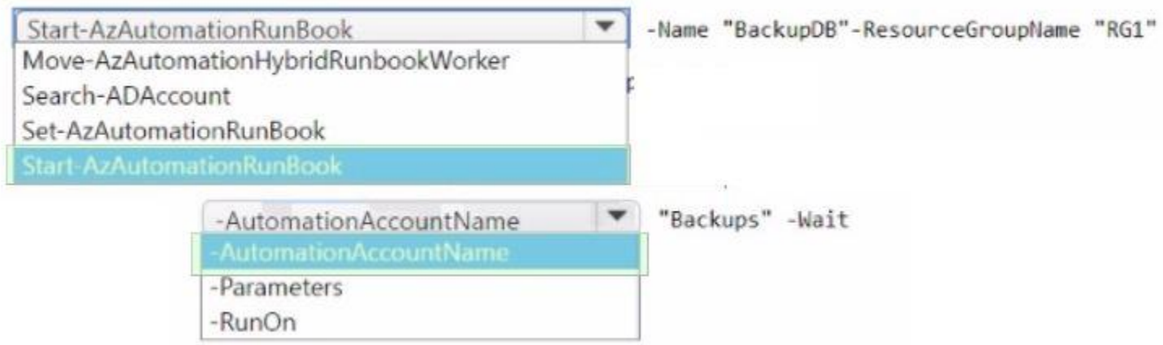

Section:
Explanation:

QUESTION 74
HOTSPOT
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a user named User1. SQLVM1 hosts a database named DB1.
You need to ensure that User1 can perform the following tasks on DB1:
* Create jobs.
* View all jobs.
* Modify, delete, and disable the jobs the user created.
The solution must use the principle of least privilege.
Which built-in database role should you assign to User1, and where is the role defined? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Answer Area:

Answer Area

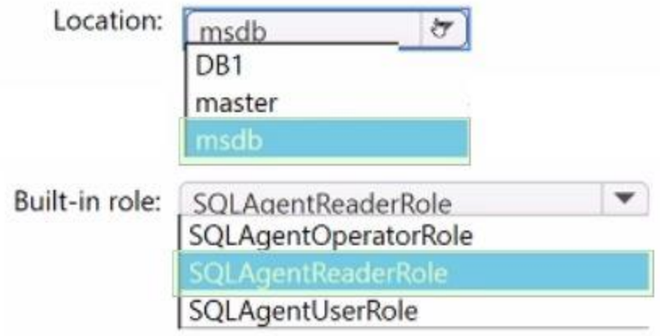

Section:
Explanation:

QUESTION 75
You plan to deploy two instances of SQL Server on Azure virtual machines in a highly available configuration that will use an Always On availability group.
You need to recommend a deployment solution that meets the following requirements:
* Provides a Service Level Agreement (SLA) of at least 99.95%
* Replicates databases in the same group synchronously
* Minimizes the latency of database writes
What should you recommend?

A. Create a proximity group and an availability set. Deploy each virtual machine to the availability set Add both virtual machines to the proximity group.
B. Create two proximity groups and a single availability set. Deploy both virtual machines to the availability set. Add one virtual machine to each proximity group.
C. Create two proximity groups and two availability sets. Deploy each virtual machine to a unique availability set. Add one virtual machine to each proximity group.
D. Create a proximity group and two availability sets. Deploy each virtual machine to a unique availability set. Add both virtual machines to the proximity group.

Correct Answer: A
Section:
Explanation:
To get VMs as close as possible, achieving the lowest possible latency, you should deploy them within a proximity placement group. https://learn.microsoft.com/en-us/azure/virtual-machines/co-location

QUESTION 76
You have an Azure subscription that contains an Azure SQL database. The database contains a table named tablet.
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 GRAIJT_FEEDBACK
C. BATCH_MODE_ADAPTIVE_TOIMS
D. BATC H_MODE_ON_ROWSTOR E

Correct Answer: D
Section:

QUESTION 77
HOTSPOT
You have an Azure subscription.
You need to deploy a logical SQL server by using PowerShell. The solution must ensure that the logical SQL server can create Azure AD users and provide Transparent Data Encryption (TDE} with a customer-managed key.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

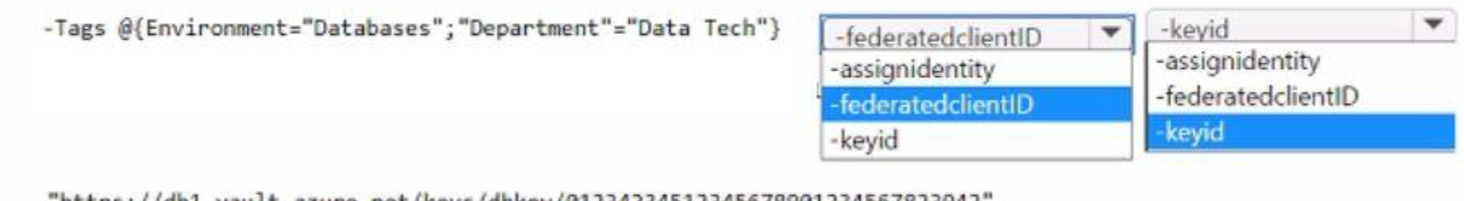

Answer Area:

Answer Area

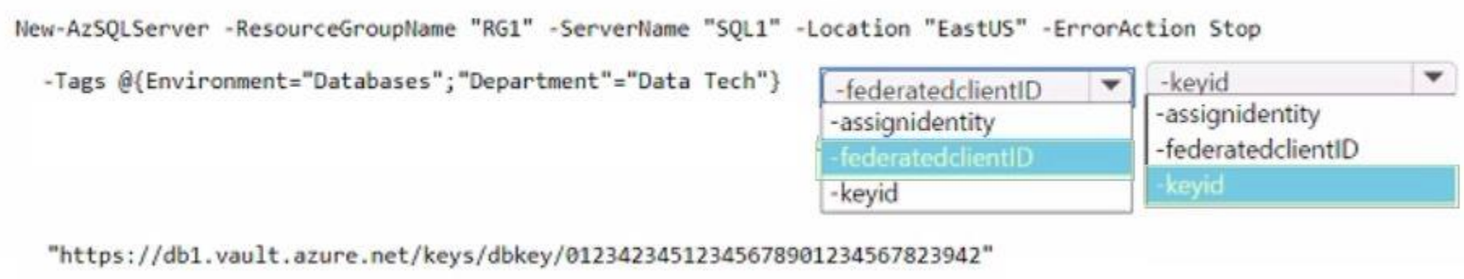

Section:
Explanation:

QUESTION 78
You have an Azure subscription.
You plan to deploy an instance of SQL Sevier on Azure Virtual Machines that supports Write Accelerator.
Which virtual machine series should you use?

A. H-series
B. G -series
C. M-series
D. E-series

Correct Answer: C
Section:

QUESTION 79
You have an Azure subscription that contains the following resources:
* 10 Azure SQL databases
* Five Azure SQL managed instances
* Five instances of SQL Server on Azure Virtual Machines
You need to implement a centralized monitoring solution for all the Azure SQL resources. The solution must minimize administrative effort. What should you include in the solution?

A. Log Analytics
B. Azure SQL Analytics
C. Query Performance Insight
D. SQL Insights

Correct Answer: B
Section:

QUESTION 80
Task 6
You need to ensure that any enhancements made to the Query Optimizer through patches are available to dbl and db2 on sql37006895.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To ensure that any enhancements made to the Query Optimizer through patches are available to dbl and db2 on sql37006895, you need to enable the query optimizer hotfixes option for each database.This option allows you to use the latest query optimization improvements that are not enabled by default1.You can enable this option by using the ALTER DATABASE SCOPED CONFIGURATION statement2.
Here are the steps to enable the query optimizer hotfixes option for dbl and db2 on sql37006895:
Connect to sql37006895 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.
Open a new query window and run the following commands for each database:
-- Switch to the database context
USE dbl;
GO
-- Enable the query optimizer hotfixes option
ALTER DATABASE SCOPED CONFIGURATION SET QUERY_OPTIMIZER_HOTFIXES = ON;
GO
Repeat the same commands for db2, replacing dbl with db2 in the USE statement.
To verify that the query optimizer hotfixes option is enabled for each database, you can query the sys.database_scoped_configurations catalog view. The value of the query_optimizer_hotfixes column should be 1 for both databases.
These are the steps to enable the query optimizer hotfixes option for dbl and db2 on sql37006895.

QUESTION 81
Task 7
You plan to create an automation runbook that will create database users in db1 from Azure AD identities. You need to configure sq1370O6895 to support the creation of new database users.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To configure sq1370O6895 to support the creation of new database users from Azure AD identities, you need to do the following steps:
Set up a Microsoft Entra tenant and associate it with your Azure subscription.You can use the Microsoft Entra portal or the Azure portal to create and manage your Microsoft Entra users and groups12.
Configure a Microsoft Entra admin for sq1370O6895.You can use the Azure portal or the Azure CLI to set a Microsoft Entra user as the admin for the server34.The Microsoft Entra admin can create other database users from Microsoft Entra identities5.
Connect to db1 using the Microsoft Entra admin account and run the following Transact-SQL statement to create a new database user from a Microsoft Entra identity:CREATE USER [Microsoft Entra user name] FROM EXTERNAL PROVIDER;6You can replace the Microsoft Entra user name with the name of the user or group that you want to create in the database.
Grant the appropriate permissions to the new database user by adding them to a database role or granting them specific privileges. For example, you can run the following Transact-SQL statement to add the new user to the db_datareader role:ALTER ROLE db_datareader ADD MEMBER [Microsoft Entra user name];
These are the steps to configure sq1370O6895 to support the creation of new database users from Azure AD identities.

QUESTION 82
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. an Azure logic app
B. a virtual machine
C. an Azure function
D. a Log Analytics workspace
E. a storage account

Correct Answer: B, D
Section:

QUESTION 83
You have an instance of SQL Server on Azure Virtual Machines named SQL1.
SQL1 contains an Extended Events session named session! 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 Set Session Event Filters settings.
B. Add a target.
C. Add an action.
D. Modify the Specify Session Data Storage settings.

Correct Answer: B
Section:

QUESTION 84
You have an on-premises datacenter that contains a 14-TB Microsoft SQL Server database.
You plan to create an Azure SQL managed instance and migrate the on-premises database to the new instance.
Which three service tiers support the SQL managed instance? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. General Purpose Standard

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

Correct Answer: A, C, D
Section:

QUESTION 85
You have an instance of SQL Server on Azure Virtual Machines named VM1.
You need to implement a disaster recovery solution that meets the following requirements:
* Returns the solution to an operational state within 15 minutes of a failure
* Can perform disaster recovery testing in an isolated environment
* Minimizes administrative effort
What should you include in the solution?

A. Azure Site Recovery
B. a failover cluster instance (FCI)
C. auto-failover groups
D. active geo-replication

Correct Answer: A
Section:
Explanation:
https://learn.microsoft.com/en-us/azure/site-recovery/site-recovery-sql#combining-bcdr-technologies-with-site-recovery

QUESTION 86
You deploy an instance of SQL Server on Azure Virtual Machines named SQL1 that hosts multiple databases.
You configure the full recovery model for all the databases.
You perform a full backup of the master database on SQL1.
You need to a perform an additional backup of the master database on SQL1. The solution must minimize how long it takes to perform the backup.
Which type of backup should you perform?

A. tail-log
B. differential
C. full
D. log

Correct Answer: B
Section:

QUESTION 87
You have an Azure SQL database named SOL1.
You need to implement a disaster recovery solution for SQL1. The solution must minimize the following:
* The recovery point objective JRPO)
* The recovery time objective (RTO)
* Administrative effort
What should you include in the solution?

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

Correct Answer: A
Section:

QUESTION 88
DRAG DROP
You have an Azure key vault named Vault1 and a SQL Server on Azure Virtual Machines instance named SQL1 SQL 1 hosts a database named DB1.
You need to configure Transparent Data Encryption (TDE) on D61 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.

Select and Place:

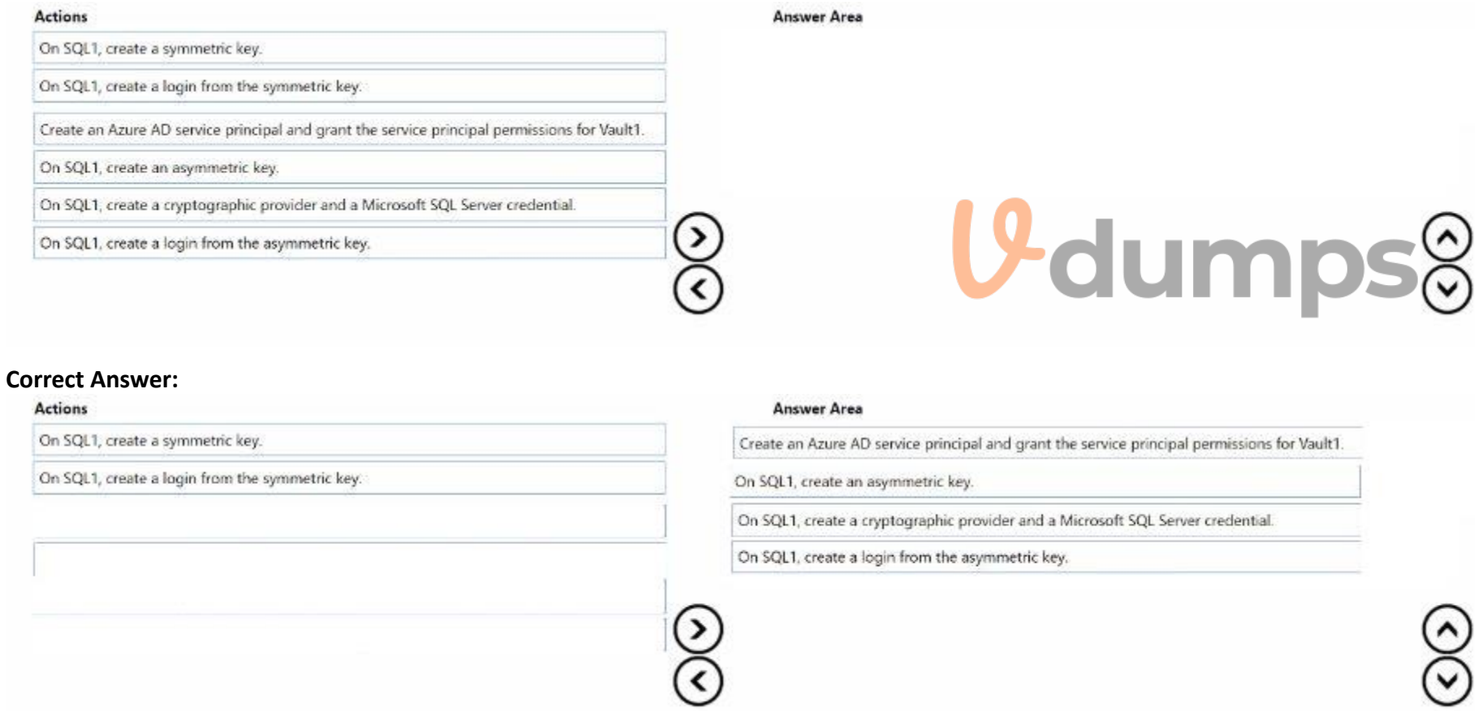

Correct Answer:

Section:
Explanation:
Create an Azure AD service principal and grant the service principal permissions for Vault1.
On SQL, 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.

QUESTION 89
You Save an Azure SCX database named DB1.

You need to query the fragmentation information of data and indexes for the tables in D61.
Which command should you run?

A.
```
sys.dm_db_index_physical_stats
```
B.
```
DBCC SHOWCONTIG
```
C.
```
DBCC CHECKALLOC
```
D.
```
sys.dm_db_index_usage_stats
```

Correct Answer: A
Section:

QUESTION 90
Task 1
In an Azure SQL database named db1, you need to enable page compression on the PK_SalesOrderHeader_SalesOrderlD clustered index of the SalesLT.SalesOrderHeader table.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To enable page compression on the PK_SalesOrderHeader_SalesOrderlD clustered index of the SalesLT.SalesOrderHeader table in db1, you can use the following Transact-SQL script:
-- Connect to the Azure SQL database named db1
USE db1;
GO
-- Enable page compression on the clustered index
ALTER INDEX PK_SalesOrderHeader_SalesOrderlD ON SalesLT.SalesOrderHeader
REBUILD WITH (DATA_COMPRESSION = PAGE);
GO
This script will rebuild the clustered index with page compression, which can reduce the storage space and improve the query performance
The script solution consists of three parts:
The first part isUSE db1; GO. This part connects to the Azure SQL database named db1, where the SalesLT.SalesOrderHeader table is located. TheGOcommand separates the batches of Transact-SQL statements and sends them to the server.
The second part isALTER INDEX PK_SalesOrderHeader_SalesOrderlD ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO. This part enables page compression on the clustered index named PK_SalesOrderHeader_SalesOrderlD, which is defined on the SalesLT.SalesOrderHeader table. TheALTER INDEXstatement modifies the properties of an existing index. TheREBUILDoption rebuilds the index from scratch, which is required to change the compression setting. TheDATA_COMPRESSION = PAGEoption specifies that page compression is applied to the index, which means that both row and prefix compression are used. Page compression can reduce the storage space and improve the query performance by compressing the data at the page level. TheGOcommand ends the batch of statements.
The third part is optional, but it can be useful to verify the compression status of the index. It isSELECT name, index_id, data_compression_desc FROM sys.indexes WHERE object_id = OBJECT_ID('SalesLT.SalesOrderHeader');. This part queries the sys.indexes catalog view, which contains information about the indexes in the database. TheSELECTstatement returns the name, index_id, and data_compression_desc columns for the indexes that belong to the SalesLT.SalesOrderHeader table. TheOBJECT_IDfunction returns the object identification number for the table name. The data_compression_desc column shows the compression type of the index, which should be PAGE for the clustered index after the script is executed.
These are the steps of the script solution for enabling page compression on the clustered index of the SalesLT.SalesOrderHeader table in db1.

QUESTION 91
Task 2
You need to configure your user account as the Azure AD admin for the server named sql3700689S.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To configure your user account as the Azure AD admin for the server named sql3700689S, you can use the Azure portal or the Azure CLI. Here are the steps for both methods:
Using the Azure portal:
Go to the Azure portal and selectSQL Server -- Azure Arc.
Select the server namedsql3700689Sand click onActive Directory admin.
Click onSet adminand choose your user account from the list of Azure AD users.
Click onSelectand thenSaveto confirm the change.
You can verify the Azure AD admin by clicking onActive Directory adminagain and checking the current admin.
Using the Azure CLI:
Install the Azure CLI and log in with your Azure account.
Run the following command to get the object ID of your user account:az ad user show --id <your-user-name> --query objectId -o tsv
Run the following command to set your user account as the Azure AD admin for the server:az sql server ad-admin create --server sql3700689S --object-id <your-object-id> --display-name <your-user-name>
You can verify the Azure AD admin by running the following command:az sql server ad-admin show --server sql3700689S
These are the steps to configure your user account as the Azure AD admin for the server named sql3700689S.

QUESTION 92
Task 3
You need to ensure that all queries executed against dbl are captured in the Query Store.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To ensure that all queries executed against dbl are captured in the Query Store, you need to enable the Query Store feature for the database and set the query capture mode to ALL.The Query Store feature provides you with insight on query plan choice and performance for Azure SQL Database1.The query capture mode controls whether all queries or only a subset of queries are tracked2.
Here are the steps to enable the Query Store and set the query capture mode to ALL for the database dbl:
Using the Azure portal:
Go to the Azure portal and select your Azure SQL Database server.
Select the database dbl and click onQuery Performance Insightin the left menu.
Click onConfigure Query Storeand turn on theQuery Storeswitch.
In theQuery Capture Modedropdown, selectAlland click onSave.
Using Transact-SQL statements:
Connect to the Azure SQL Database server and the database dbl using SQL Server Management Studio or Azure Data Studio.
Run the following command to enable the Query Store for the database:ALTER DATABASE dbl SET QUERY_STORE = ON;
Run the following command to set the query capture mode to ALL for the database:ALTER DATABASE dbl SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL);
These are the steps to ensure that all queries executed against dbl are captured in the Query Store.

QUESTION 93
Task 4
You need to enable change data capture (CDC) for db1.

A. See the explanation part for the complete Solution

Correct Answer: A

Section:
Explanation:
To enable change data capture (CDC) for db1, you need to run the stored procedure sys.sp_cdc_enable_db in the database context.CDC is a feature that records activity on a database when tables and rows have been modified1.CDC can be used for various scenarios, such as data synchronization, auditing, or ETL processes2.
Here are the steps to enable CDC for db1:
Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.
Open a new query window and run the following command:EXEC sys.sp_cdc_enable_db; GO
This command will enable CDC for the database and create the cdc schema, cdc user, metadata tables, and other system objects for the database3.
To verify that CDC is enabled for db1, you can query the is_cdc_enabled column in the sys.databases catalog view. The value should be 1 for db1.
These are the steps to enable CDC for db1

QUESTION 94
Task 5
You need to configure a disaster recovery solution for db1. When a failover occurs, the connection strings to the database must remain the same. The secondary server must be in the West US 3 Azure region.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To configure a disaster recovery solution for db1, you can use the failover groups feature of Azure SQL Database.Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings1.You can also use active geo-replication as an alternative, but you will need to update the connection strings manually after a failover2.
Here are the steps to create a failover group for db1 with the secondary server in the West US 3 region:
Using the Azure portal:
Go to the Azure portal and select your Azure SQL Database server that hosts db1.
SelectFailover groupsin the left menu and click onAdd group.
Enter a name for the failover group and selectWest US 3as the secondary region.
Click onCreate a new serverand enter the details for the secondary server, such as server name, admin login, password, and subscription.
Click onSelect existing database(s)and choose db1 from the list of databases on the primary server.
Click onConfigure failover policyand select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.
Click onCreateto create the failover group and start the replication of db1 to the secondary server.
Using PowerShell commands:
Install the Azure PowerShell module and log in with your Azure account.
Run the following command to create a new server in the West US 3 region:New-AzSqlServer -ResourceGroupName <your-resource-group-name> -ServerName <your-secondary-server-name> -Location 'West US 3' -SqlAdministratorCredentials $(New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList '<your-admin-login>', $(ConvertTo-SecureString -String '<your-password>' -AsPlainText -Force))
Run the following command to create a new failover group with db1:New-AzSqlDatabaseFailoverGroup -ResourceGroupName <your-resource-group-name> -ServerName <your-primary-server-name> -PartnerResourceGroupName <your-resource-group-name> -PartnerServerName <your-secondary-server-name> -FailoverGroupName <your-failover-group-name> -Database db1 -FailoverPolicy Manual -GracePeriodWithDataLossHours 1 -ReadWriteFailoverEndpoint 'Enabled'
You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.
These are the steps to create a failover group for db1 with the secondary server in the West US 3 region.

QUESTION 95
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 laaS Agent extension (SqllaasExtension). 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

Correct Answer: A
Section:

QUESTION 96
DRAG DROP
You have an Azure subscription that contains an instance of SQL Server on Azure Virtual Machines named SQLVM1 and a virtual machine named Server! that runs Windows Server. SQLVM1 and Serverl are joined to an Active Directory Domain Services (AD DS) domain. Serverl hosts a file share named Sharel.
You need to ensure that a SQL Server Agent job step on SQLVM1 can access the files in Share1. The solution must use the principle of least privilege.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

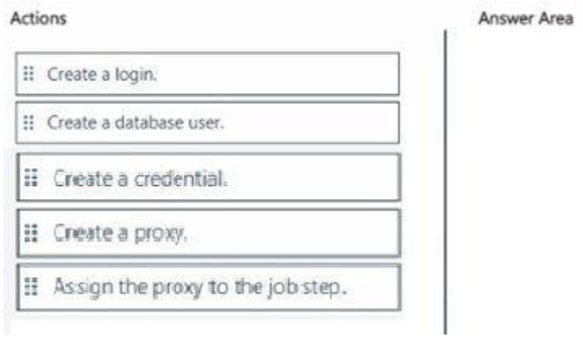

Correct Answer:

Section:
Explanation:

QUESTION 97
HOTSPOT
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 Server! 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.

Hot Area:

Answer Area

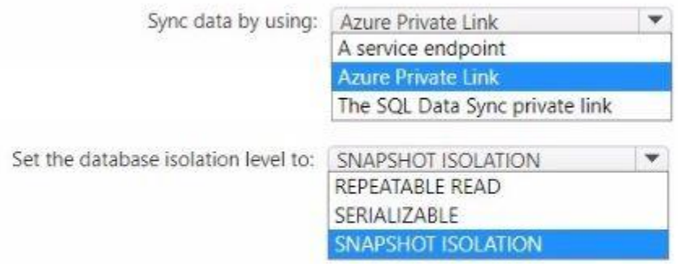

Answer Area:

Answer Area

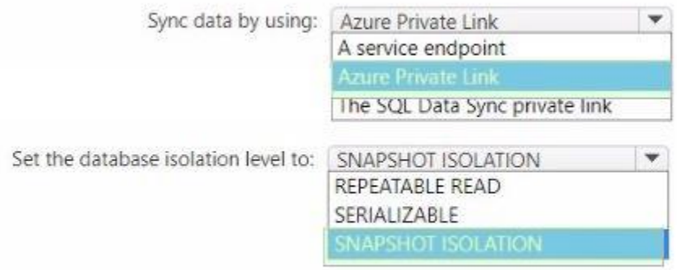

Section:
Explanation:

QUESTION 98
You have an Azure virtual machine named VM1 that runs Windows Server 2022 and hosts a Microsoft SQL Server 2019 instance named SQL1.
You need to configure SQL1 to use mixed mode authentication.
Which procedure should you run?

A. sp_eddremotelogin
B. xp_grant_login
C. sp_change_users_login
D. xp_instance_regwrite

Correct Answer: D
Section: