Number: DP-300
Passing Score: 800
Time Limit: 120
File Version: 4.0

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 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:

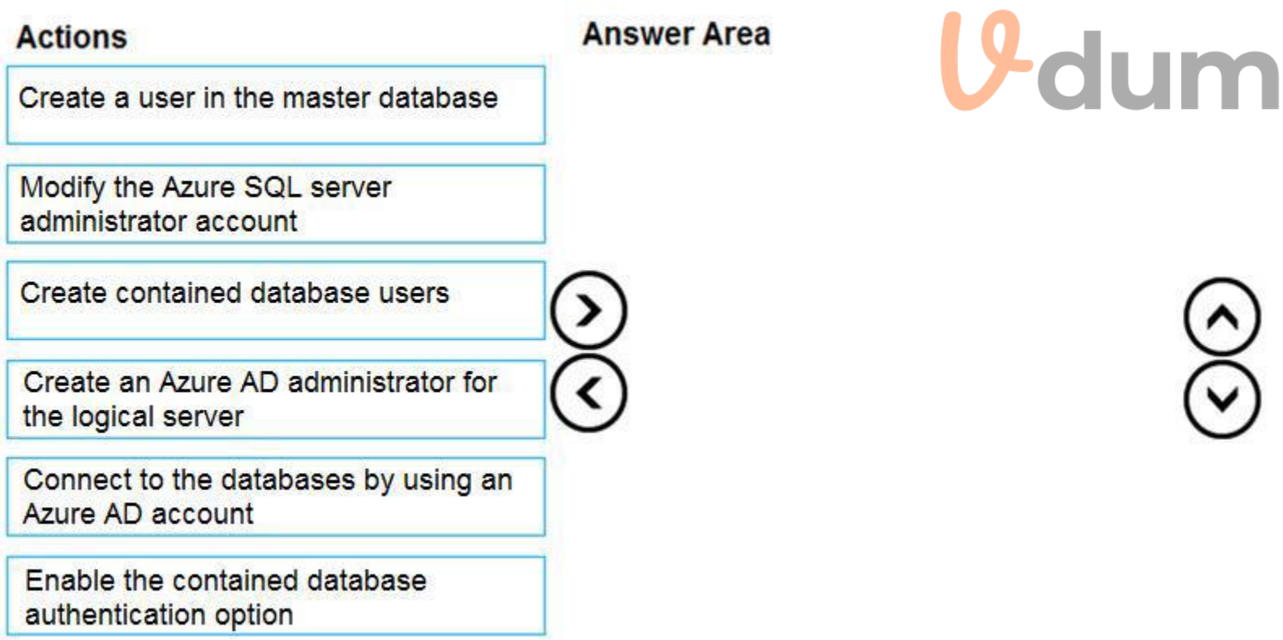

Correct Answer:

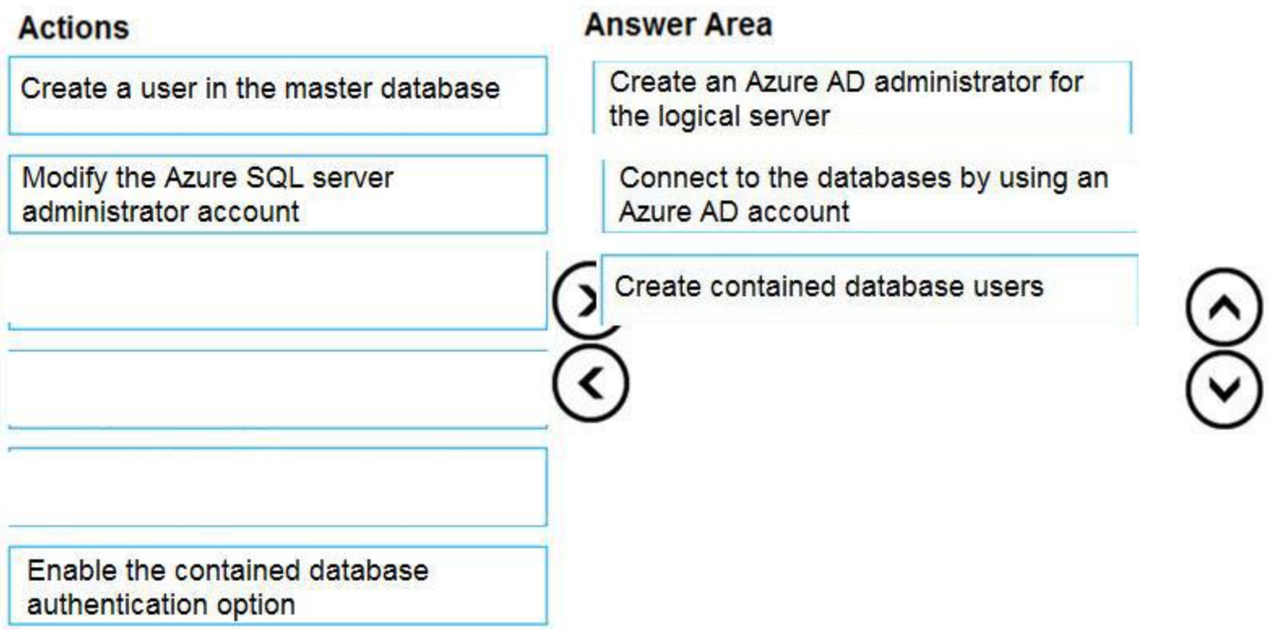

Section:
Explanation:
Scenario: Authenticate database users by using Active Directory credentials.
The configuration steps include the following procedures to configure and use Azure Active Directory authentication.
1. Create and populate Azure AD.
2. Optional: Associate or change the active directory 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

QUESTION 2
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:

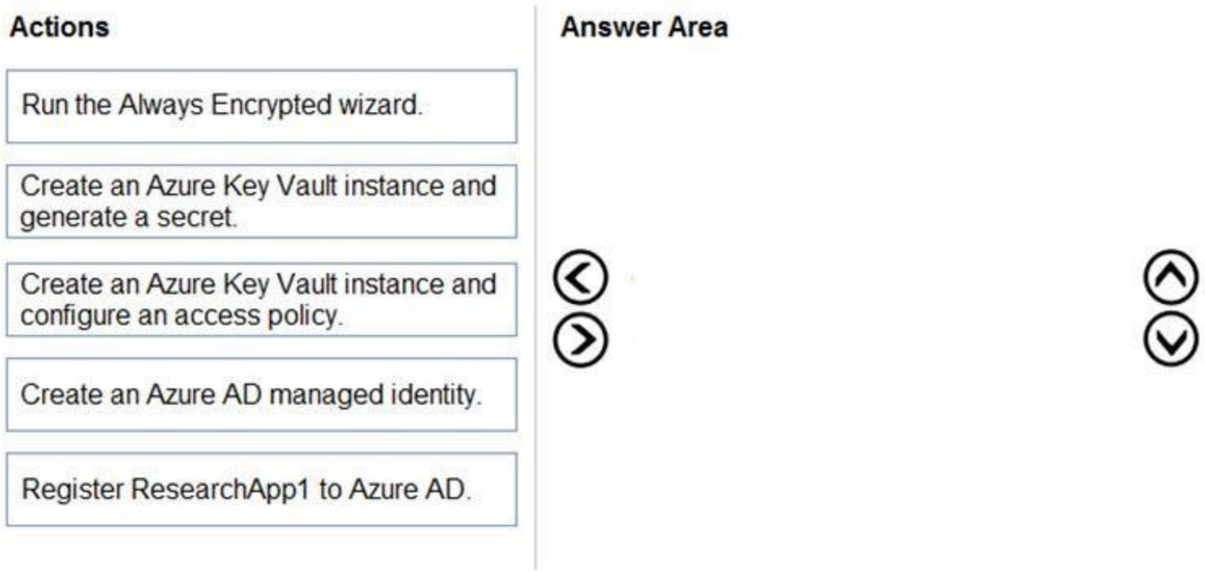

Correct Answer:

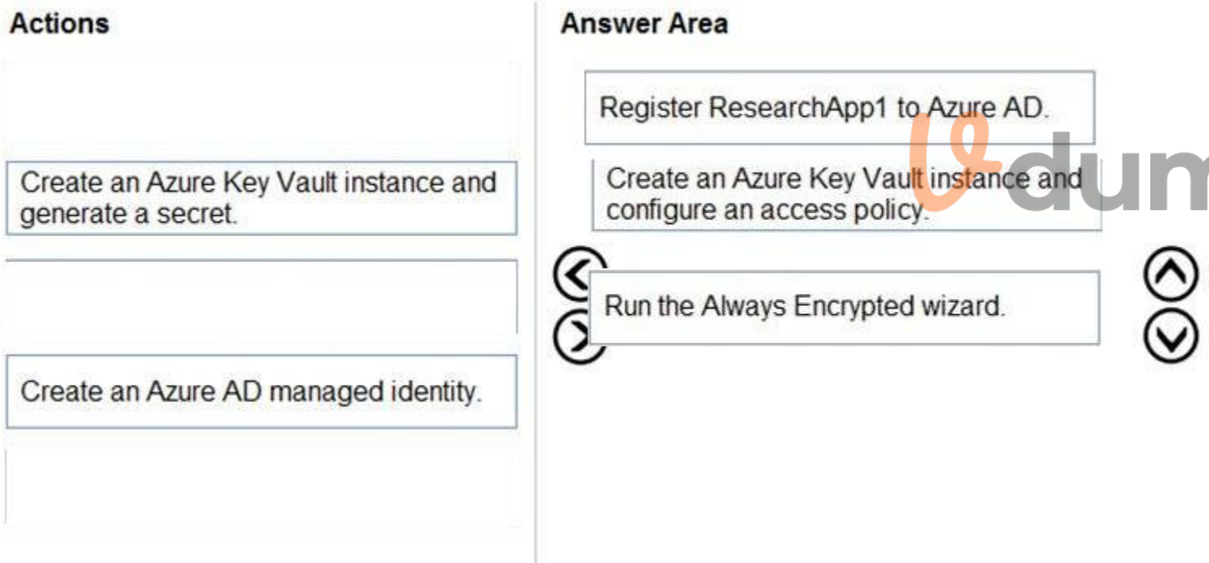

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/always-encrypted-azure-key-vault-configure?tabs=azure-powershell

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.

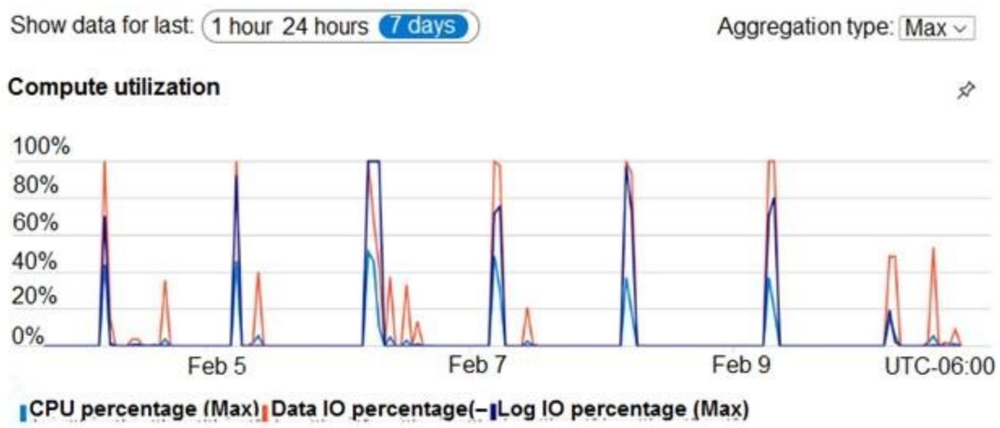

Role Assignments

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

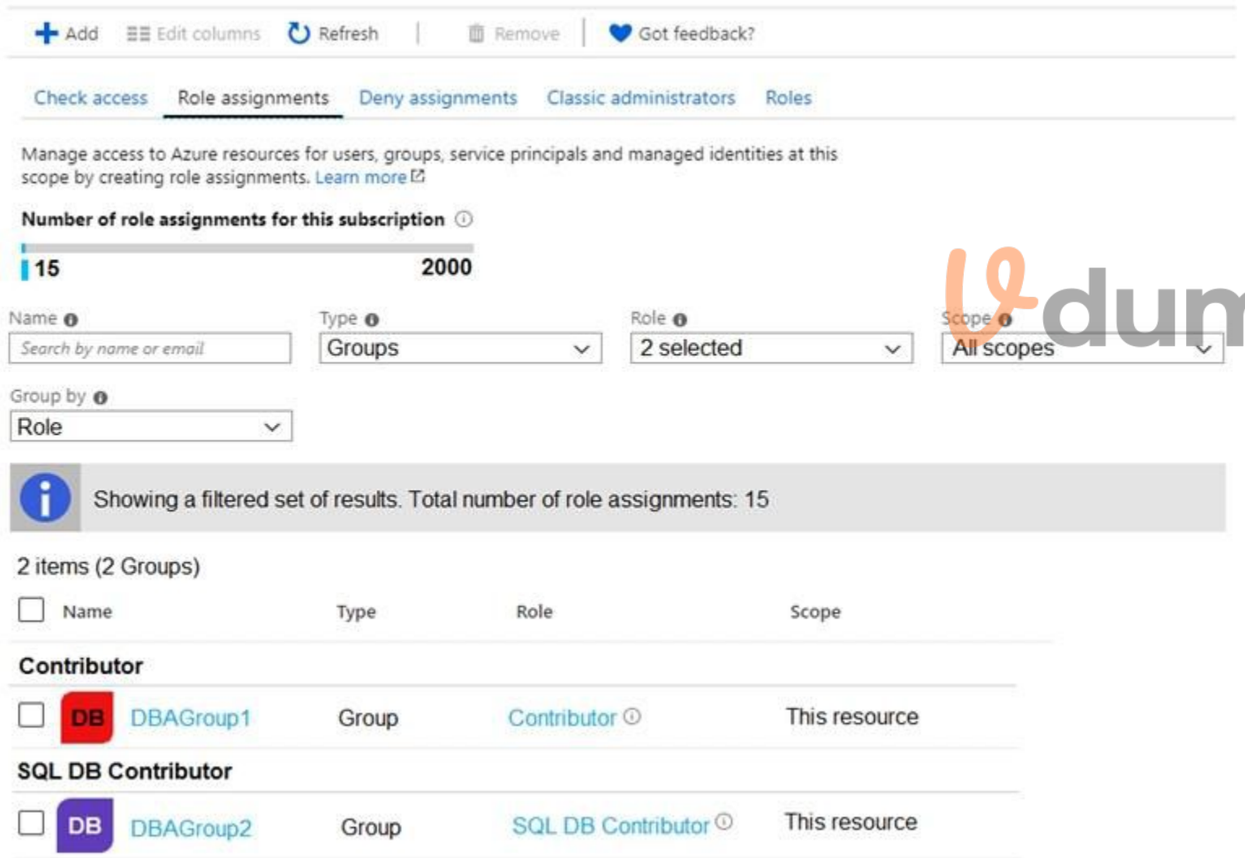

QUESTION 1

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

Statements	Yes	No
DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.	○	○
DBAGroup1 will be able to assign the SQL DB Contributor role to other users.	○	○
DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server.	○	○

Answer Area:

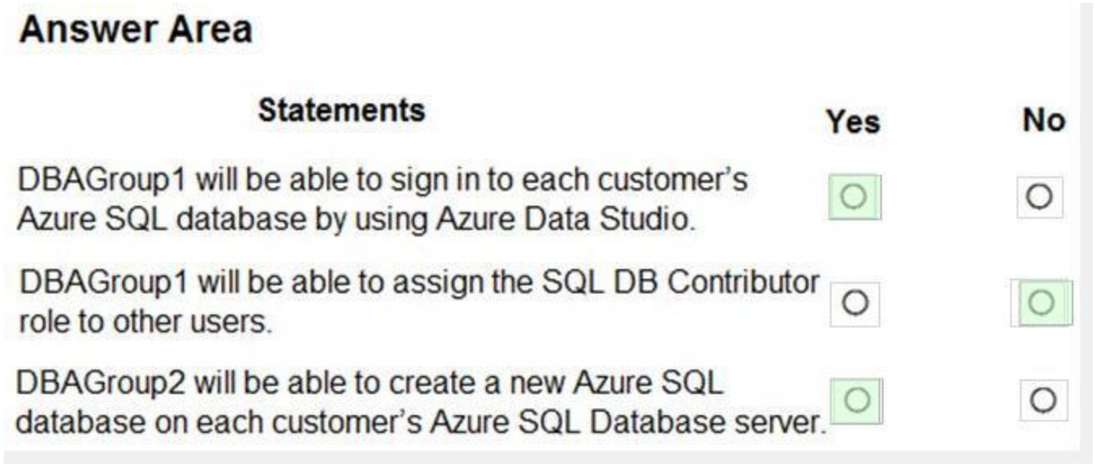

Section:
Explanation:
Box 1: Yes
DBAGroup1 is member of the Contributor role.
The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.
Box 2: No
Box 3: Yes
DBAGroup2 is member 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
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

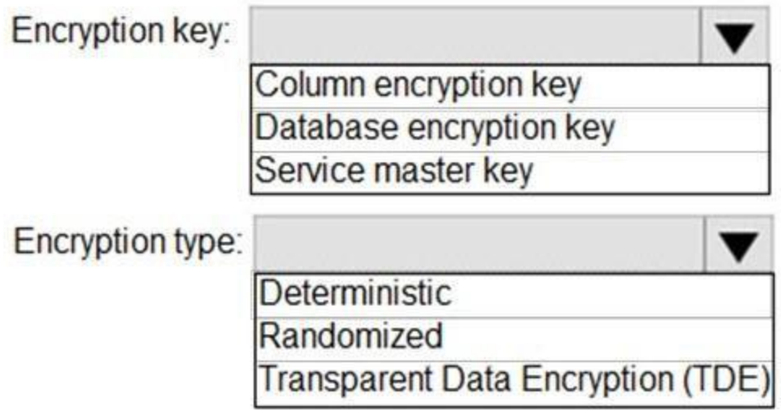

Answer Area:

Answer Area

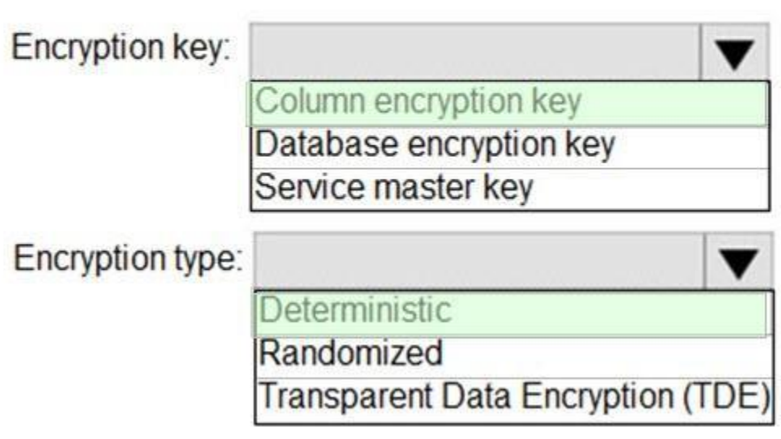

Section:

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 2
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:

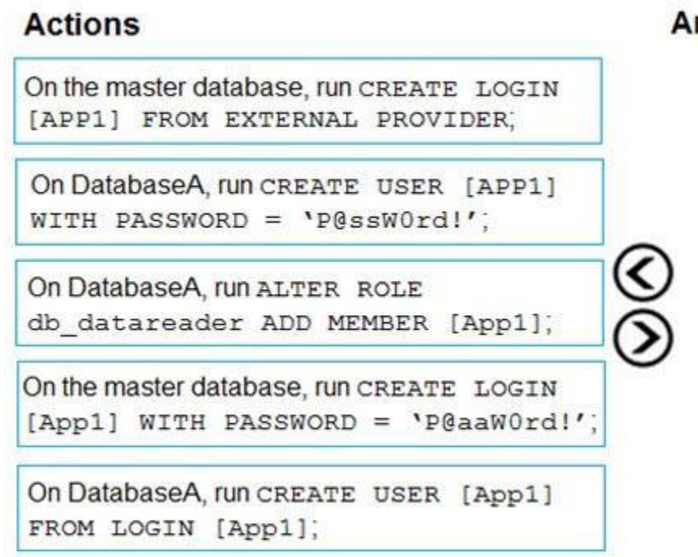

Correct Answer:

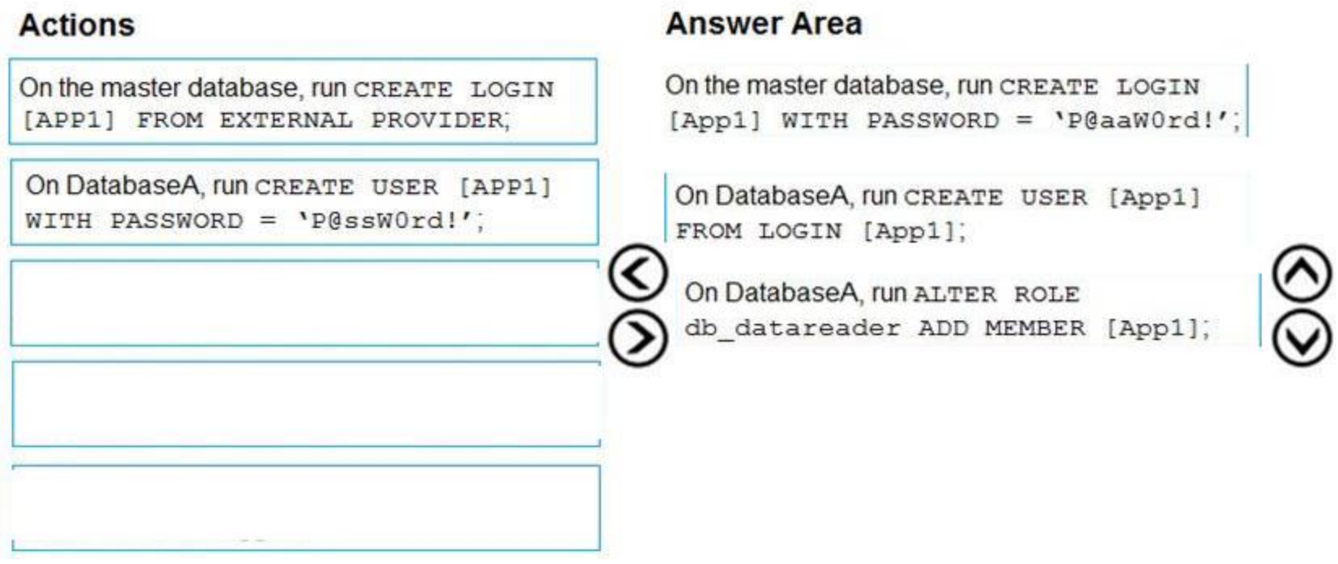

Section:
Explanation:
Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'
Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:
CREATE LOGIN readonlylogin WITH 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 3
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 4
You are designing an enterprise data warehouse in Azure Synapse Analytics that will contain a table named Customers. Customers will contain credit card information. You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers. The solution must prevent all the salespeople from viewing or inferring the credit card information. What should you include in the recommendation?

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

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 5
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 6
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 7
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 8
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 9
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 10
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.

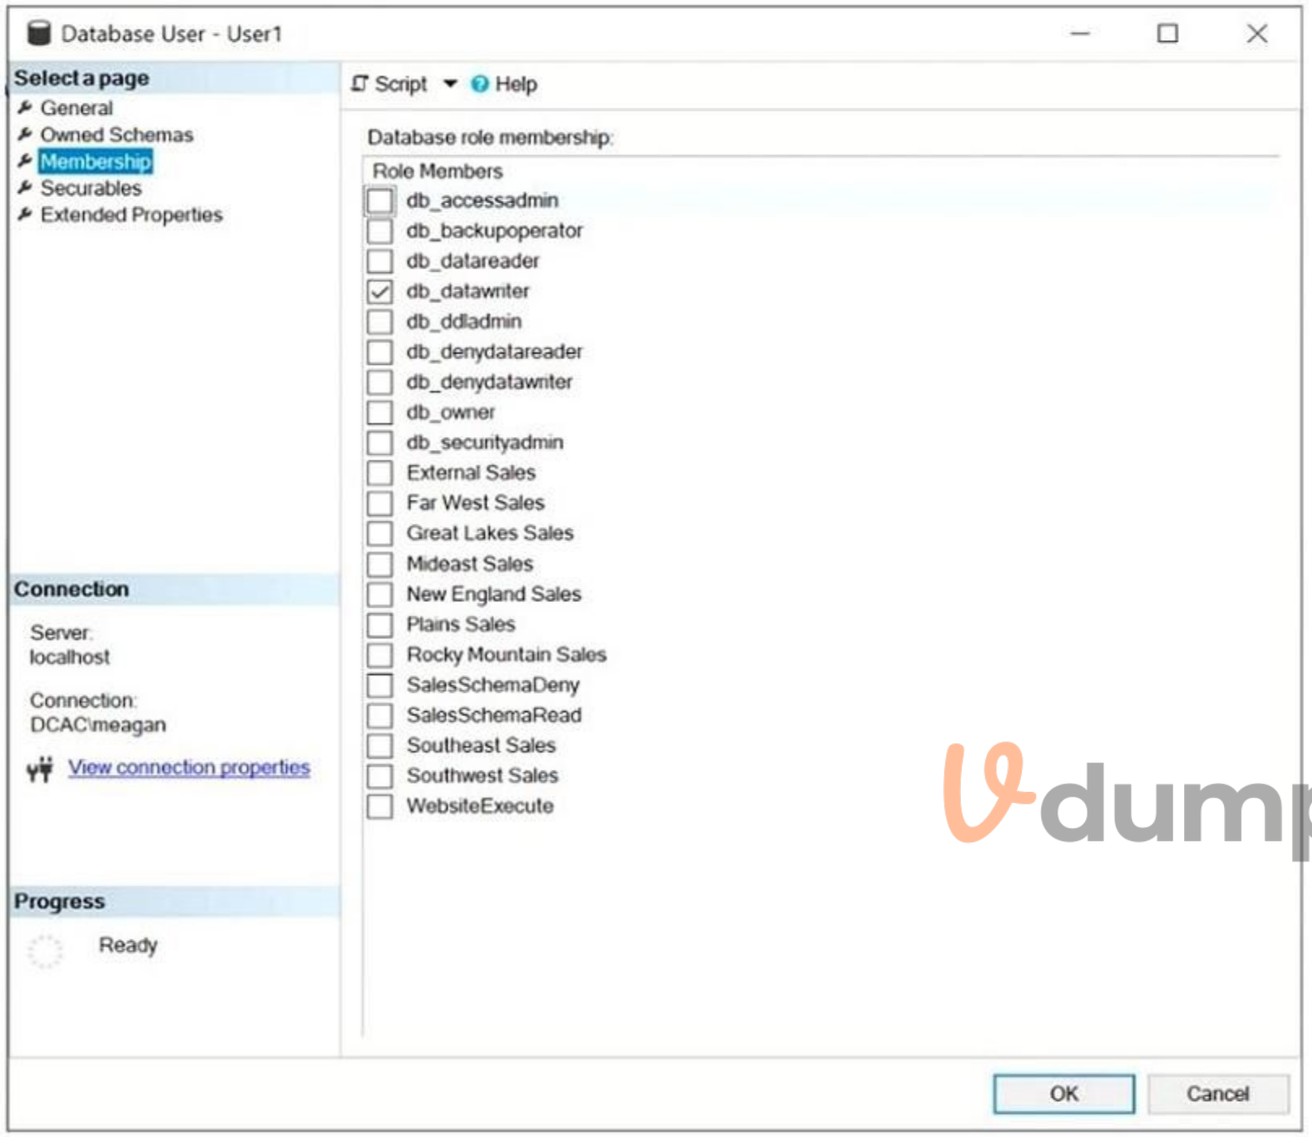

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

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 11
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:

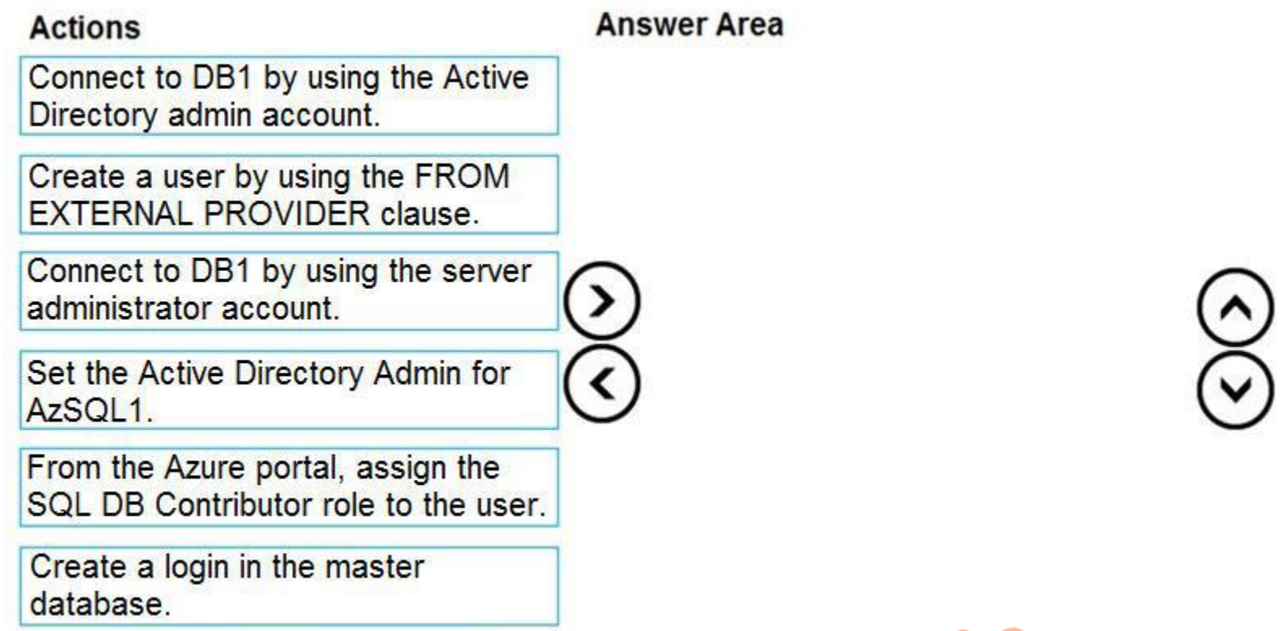

Correct Answer:

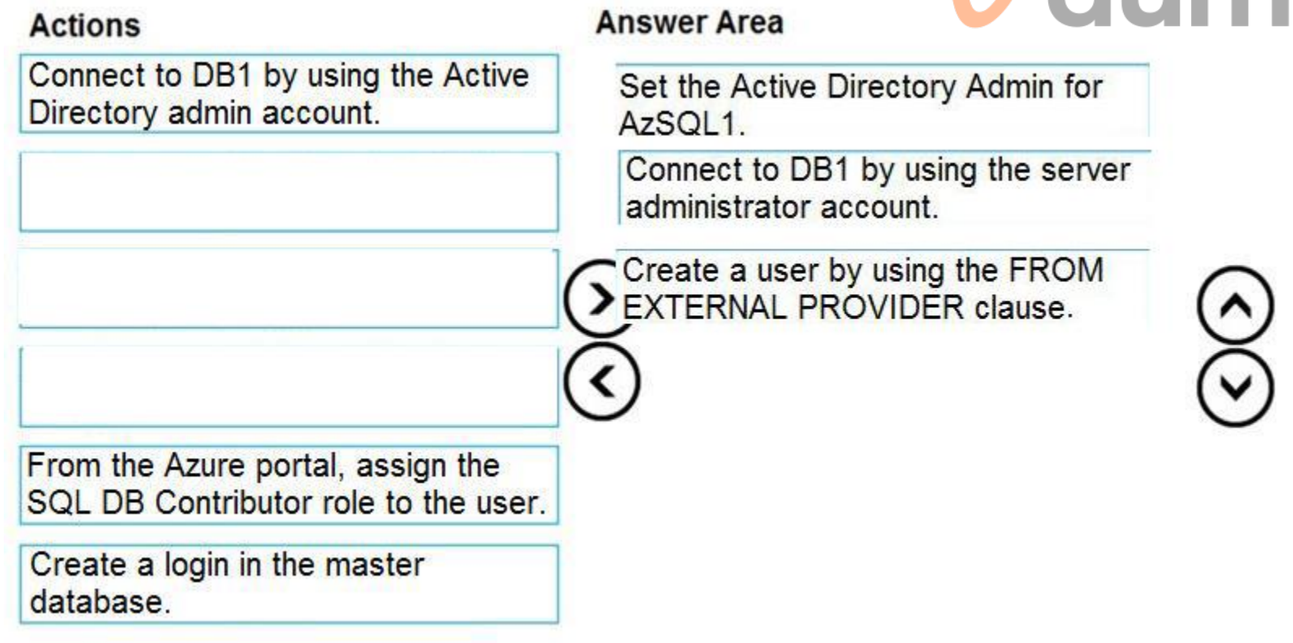

Section:
Explanation:
Step 1: Set up the Active Directory Admin for AzSQL1.
Step 2: Connect to DB1 by using the server administrator.

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 12
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

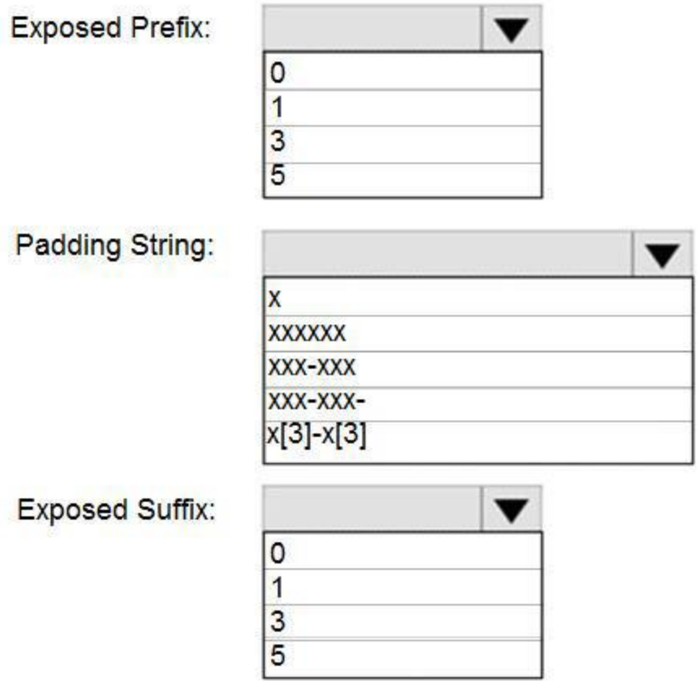

Answer Area:

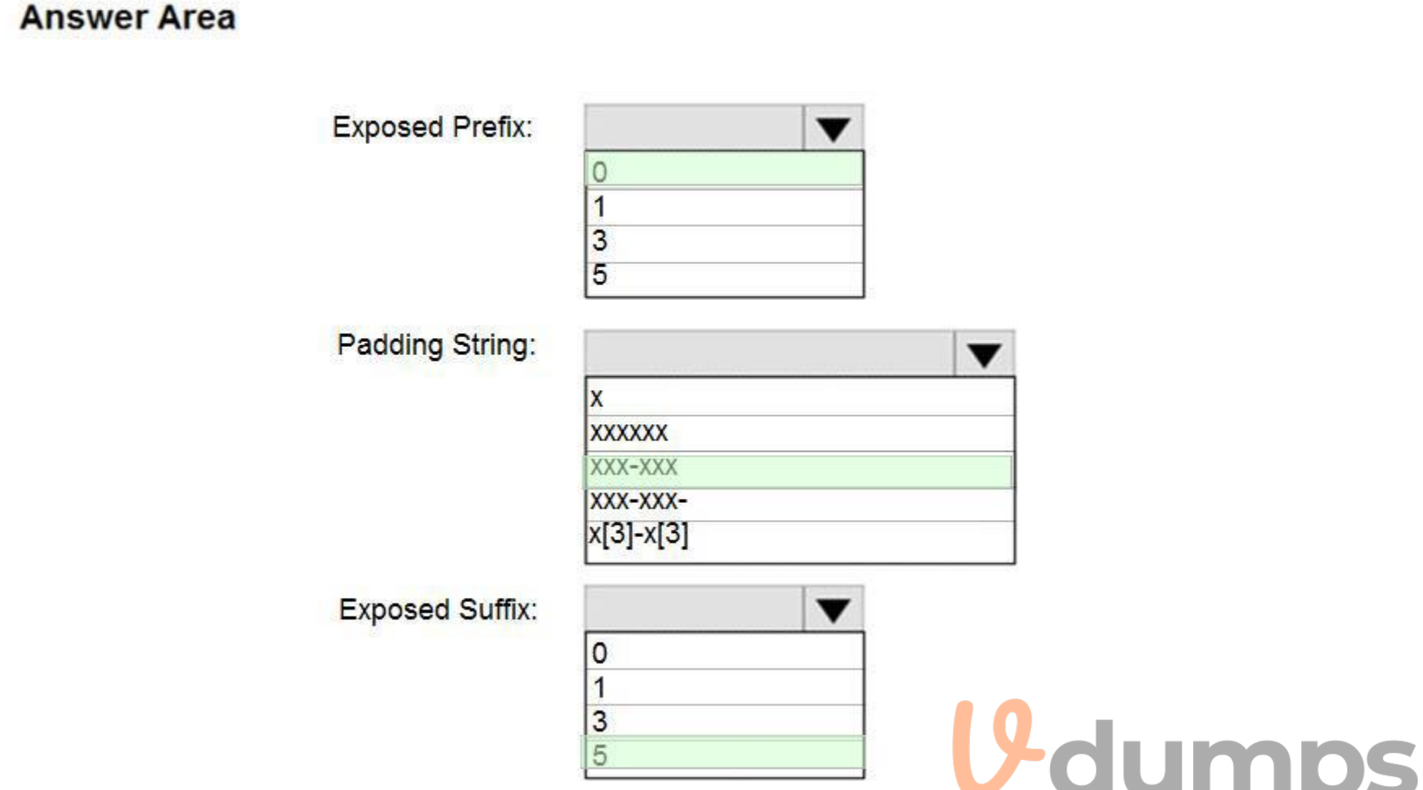

Section:

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 13
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:

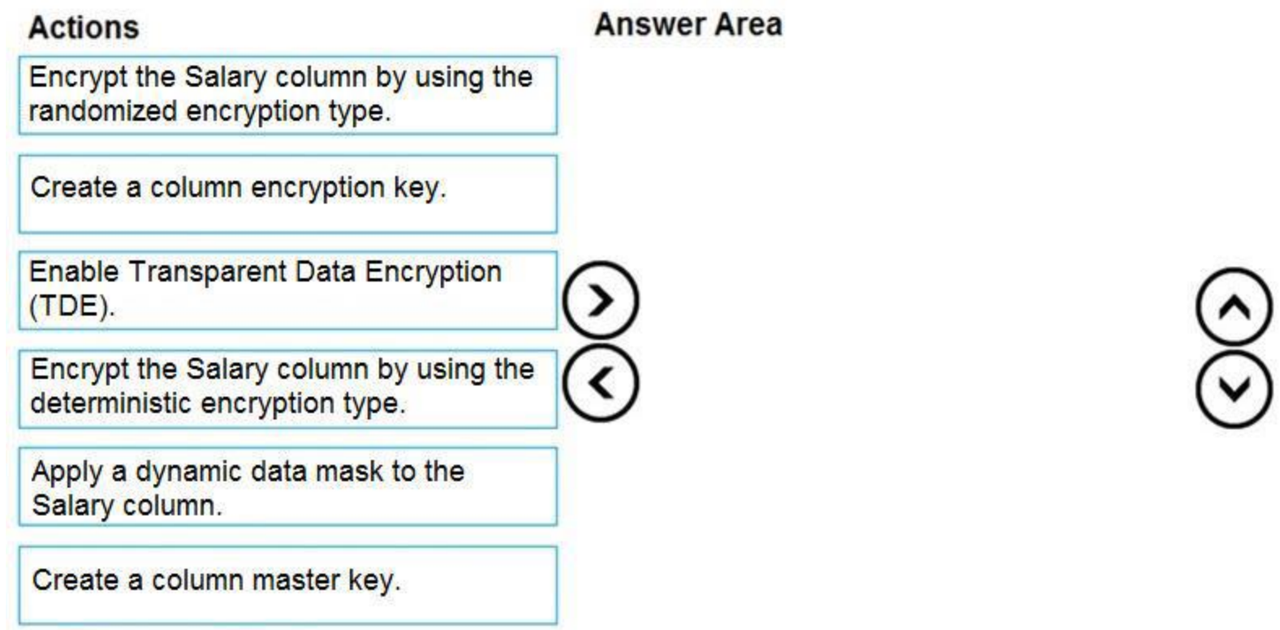

Correct Answer:

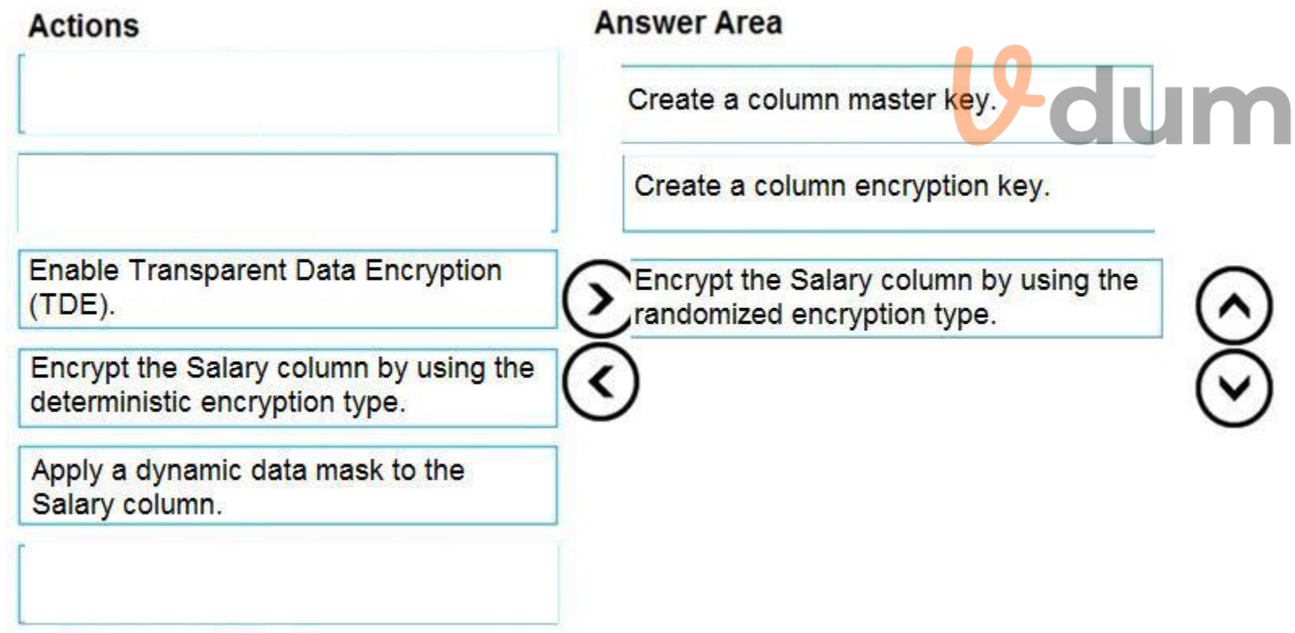

Section:

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 14
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 15
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 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:

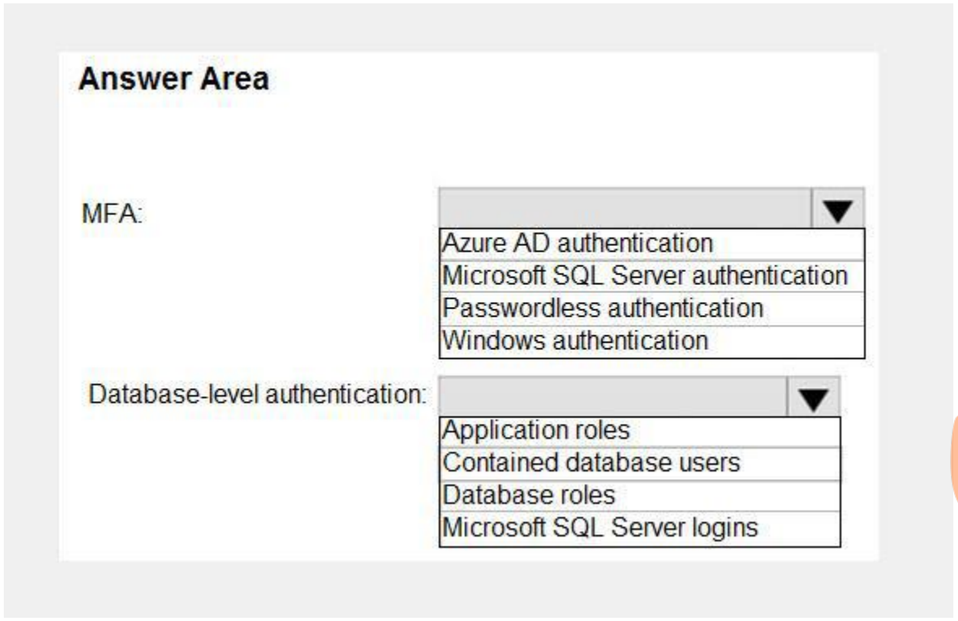

Answer Area:

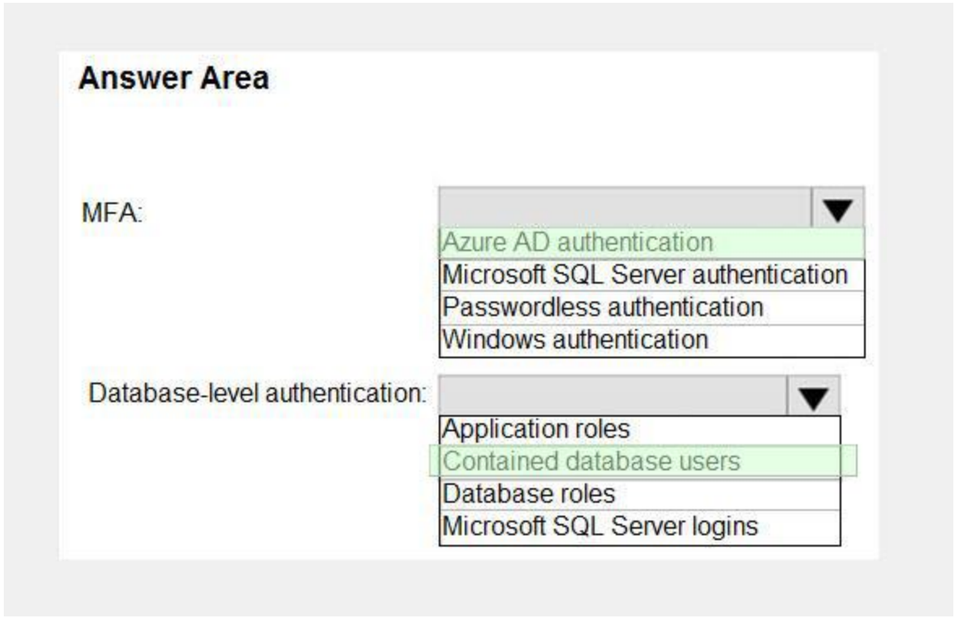

Section:
Explanation:
Box 1: Azure AD authentication
Azure Active Directory authentication supports Multi-Factor authentication through Active Directory Universal Authentication.
Box 2: Contained database users
Azure Active Directory Uses contained database users to authenticate identities at the database 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 - Monitor and Optimize Operational Resources
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 implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.
How should you collect and stream metrics? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

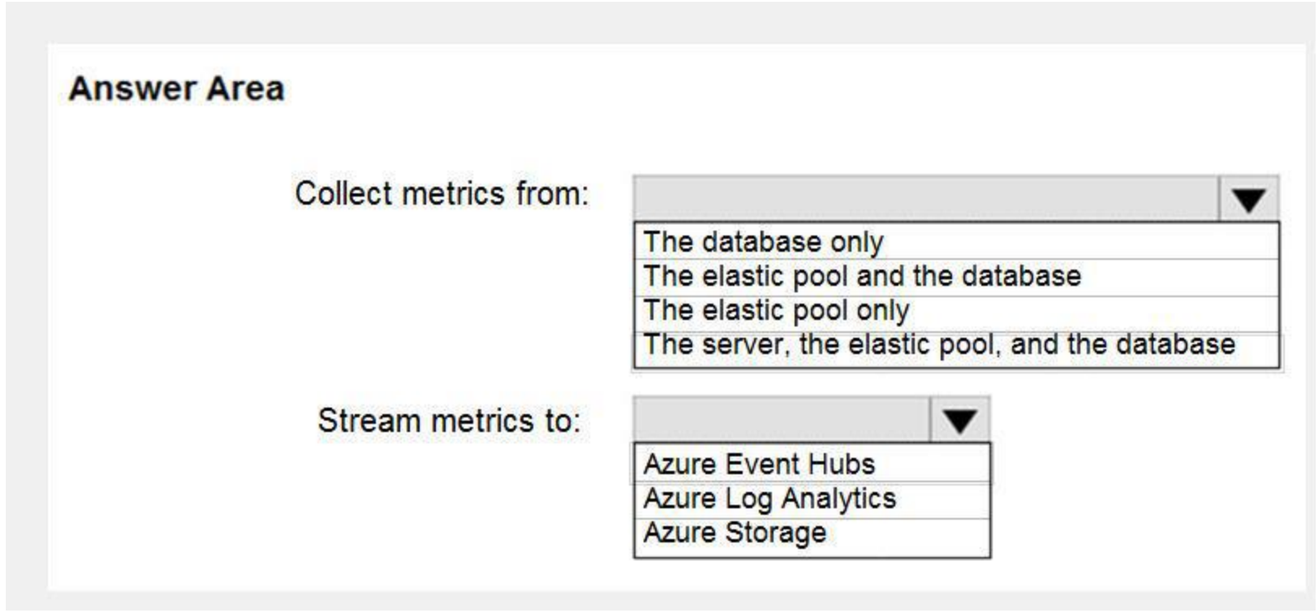

Answer Area:

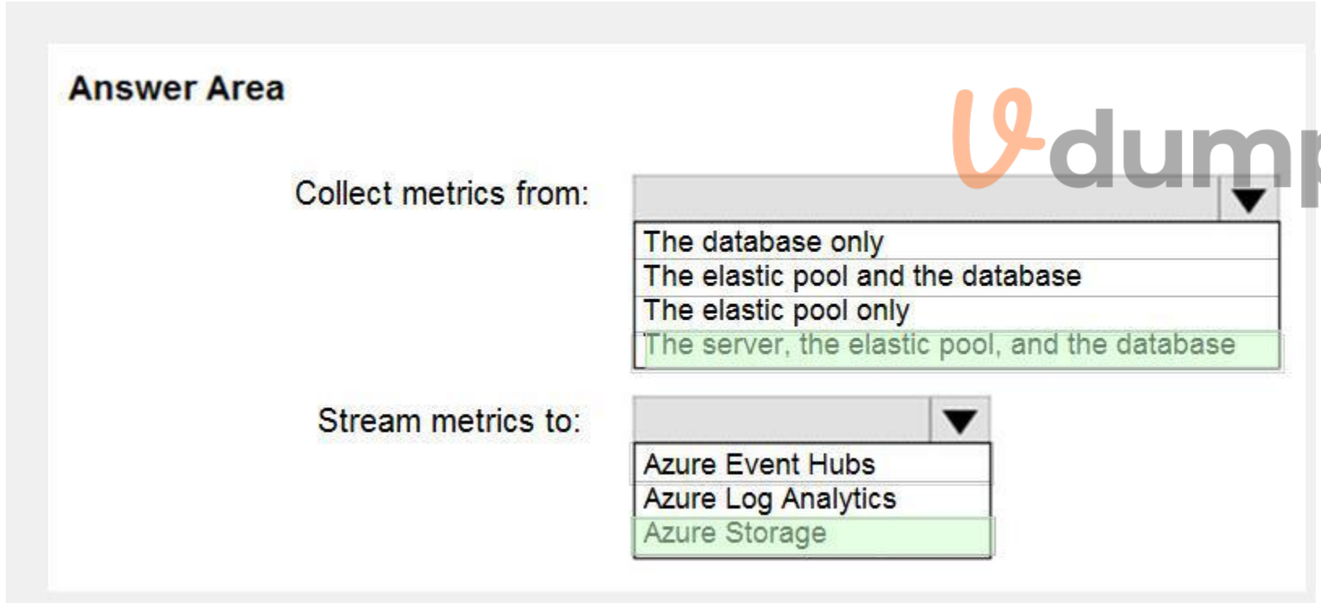

Section:
Explanation:
Box 1: The server, the elastic pool, and the database
Senario:
SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.
Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality. Box 2: Azure Event hubs
Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
Event hubs are able to handle custom metrics.
Incorrect Answers:
Azure Log Analytics

Azure metric and log data are sent to Azure Monitor Logs, previously known as Azure Log Analytics, directly by Azure. Azure SQL Analytics is a cloud only monitoring solution supporting streaming of diagnostics telemetry for all of your Azure SQL databases.
However, because Azure SQL Analytics does not use agents to connect to Azure Monitor, it does not support monitoring of SQL Server hosted on-premises or in virtual machines.

02 - Monitor and Optimize Operational Resources
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.

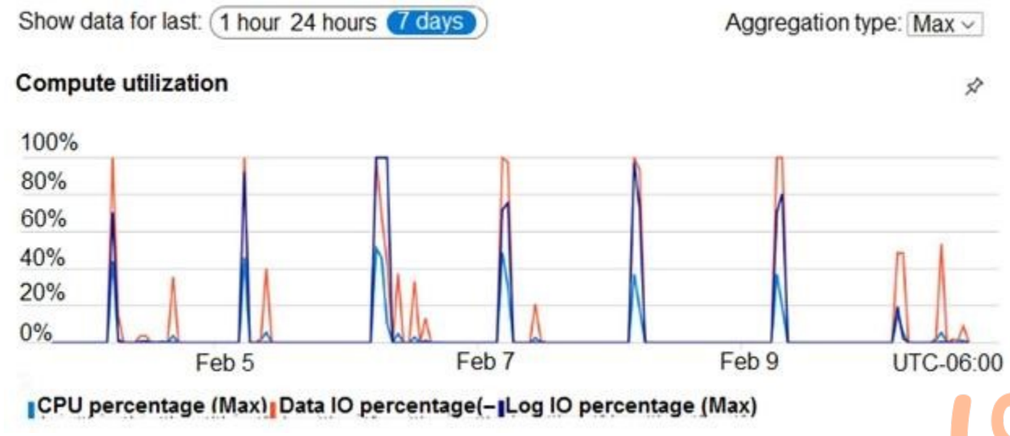

Role Assignments
For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

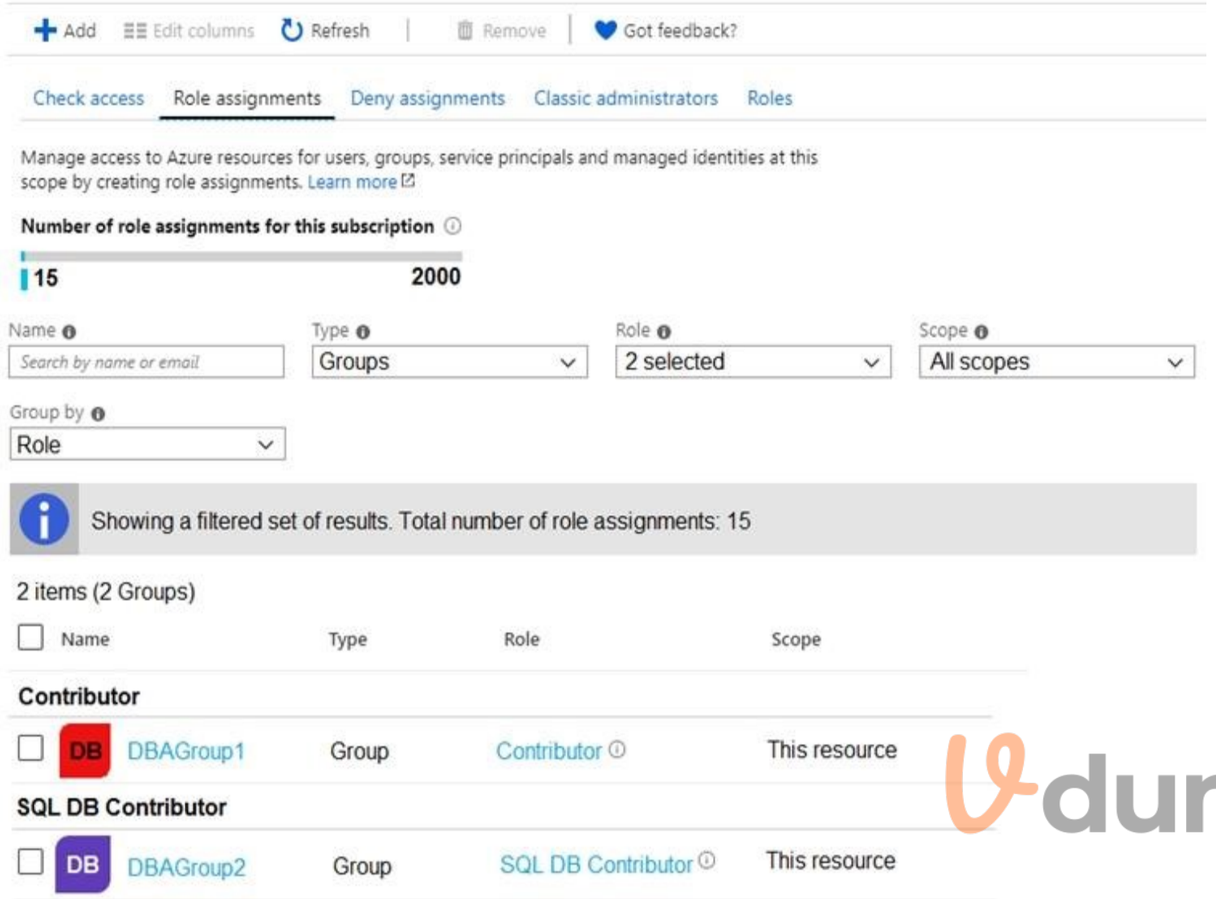

QUESTION 1
Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

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

Correct Answer: A
Section:
Explanation:
There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed.
Incorrect Answers:
B: Hyperscale is for large databases
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

QUESTION 2
Which audit log destination should you use to meet the monitoring requirements?

A. Azure Storage
B. Azure Event Hubs
C. Azure Log Analytics

Correct Answer: C
Section:
Explanation:
Scenario: Use a single dashboard to review security and audit data for all the PaaS databases.
With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics. Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards

03 - Monitor and Optimize Operational Resources
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
ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.
Existing Environment
ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.
SALESDB collects data from the stores and the website.
DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.
REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.
Requirements
Planned Changes
ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:
Migrate SALESDB and REPORTINGDB to an Azure SQL database.
Migrate DOCDB to Azure Cosmos DB.
The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.
As they arrive, all the sales documents in JSON format must be transformed into one consistent format. Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.
Technical Requirements
The new Azure data infrastructure must meet the following technical requirements:
Data in SALESDB must encrypted by using Transparent Data Encryption (TDE). The encryption must use your own key. SALESDB must be restorable to any given minute within the past three weeks.
Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns. Missing indexes must be created automatically for REPORTINGDB.
Disk IO, CPU, and memory usage must be monitored for SALESDB.

QUESTION 1
Which windowing function should you use to perform the streaming aggregation of the sales data?

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

Correct Answer: D
Section:
Explanation:
Scenario: The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.
Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

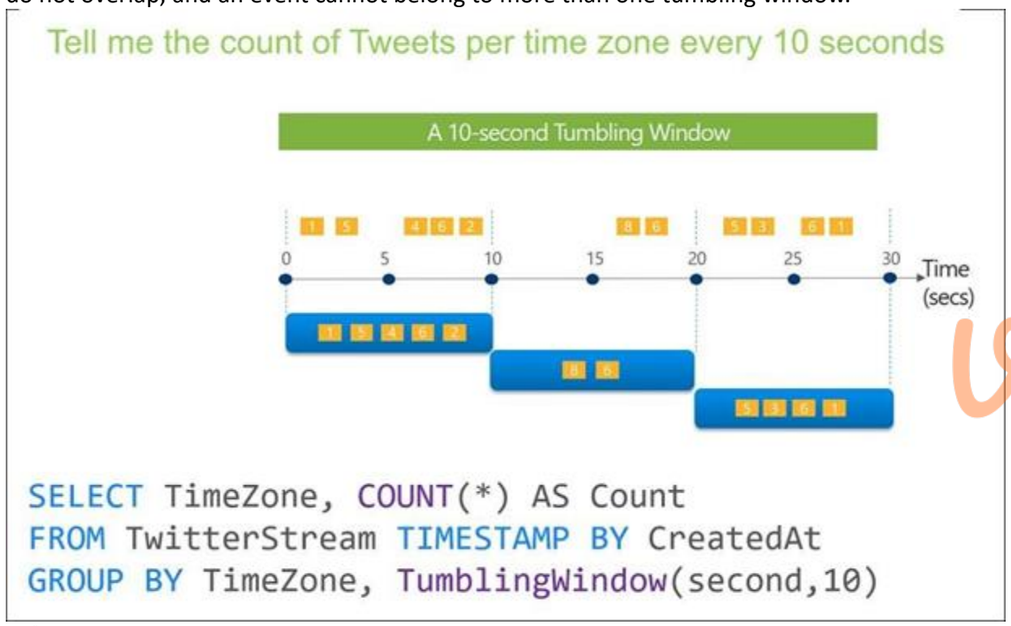

Reference:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-functions.md

QUESTION 2
Which counter should you monitor for real-time processing to meet the technical requirements?

A. SU% Utilization
B. CPU% utilization
C. Concurrent users
D. Data Conversion Errors

Correct Answer: B
Section:
Explanation:
Scenario: Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns. To monitor the performance of a database in Azure SQL Database and Azure SQL Managed Instance, start by monitoring the CPU and IO resources used by your workload relative to the level of database performance you chose in selecting a particular service tier and performance level.
Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/monitor-tune-overview Question Set 4

QUESTION 3
HOTSPOT
You plan to deploy Instance1 by using the following script.

```
{
    "type": "Microsoft.Sql/managedInstances",
    "apiVersion": "2019-06-01-preview",
    "name": "[parameters('instanceName')]",
    "location": "[resourceGroup().location]",
    "sku": {
        "name": "[parameters('skuName')]",
        "tier": "[parameters('skuEdition')]"
    },
    "dependsOn": [
        "Microsoft.Resources/deployments/BuildMINetworking"
    ],
    "properties": {
        "administratorLogin": "[parameters('adminLogin')]",
        "administratorLoginPassword0": "[parameters('adminPassword')]",
        "subnetId": "
[resourceId('Microsoft.Network/virtualNetworks/subnets',parameters('netName'),parameters('subnetName'))]",
```

You need to specify the licenseType and storagenedundancy parameters. The deployment must meet the availability requirements and the business requirements for DB1 and DB2.
To what should you set each parameter? To answer, select the appropriate options in the answer area.

Hot Area:

Answer Area:

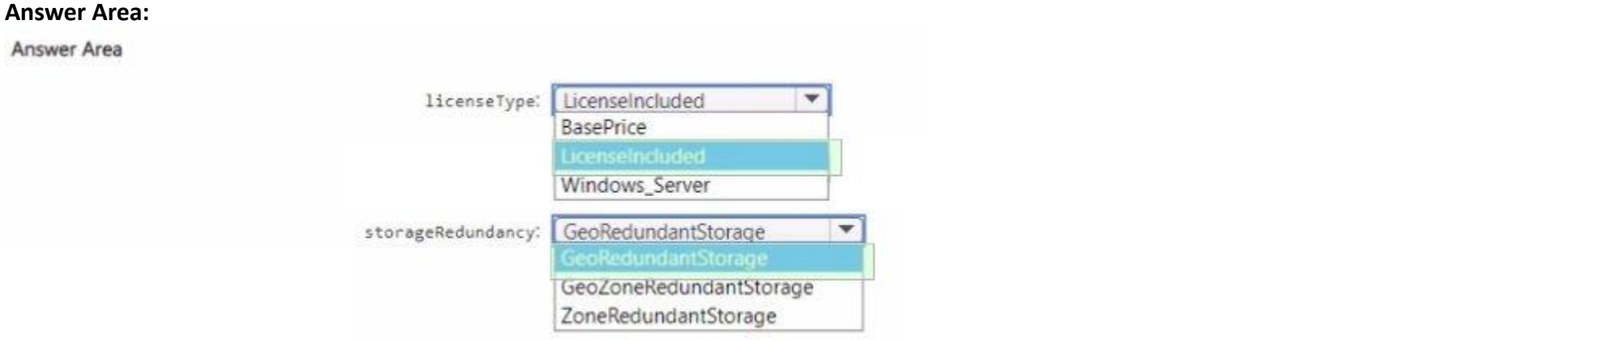

Section:
Explanation:

QUESTION 4
You need to recommend a backup solution to restore DB3. The solution must meet the availability requirements. Which type of backup should you use?

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

Correct Answer: C
Section:

QUESTION 5
You need to recommend which configuration to perform twice to enable access to the primary and secondary replicas of DB3. The solution must meet the availability requirements.
What should you recommend?

A. Configure virtual network service endpoints.
B. Enable database firewall rules.
C. Create database-scoped credentials.
D. Configure connection strings that reference the read-write listener.

Correct Answer: D
Section:

QUESTION 6
DRAG DROP
You need to recommend an authentication solution for App1 access to DB1 and DB2 after their migration to Instance1. The solution must meet the availability requirements.
Which actions should you perform in sequence? To answer, drag the appropriate actions to the correct order. Each action 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:

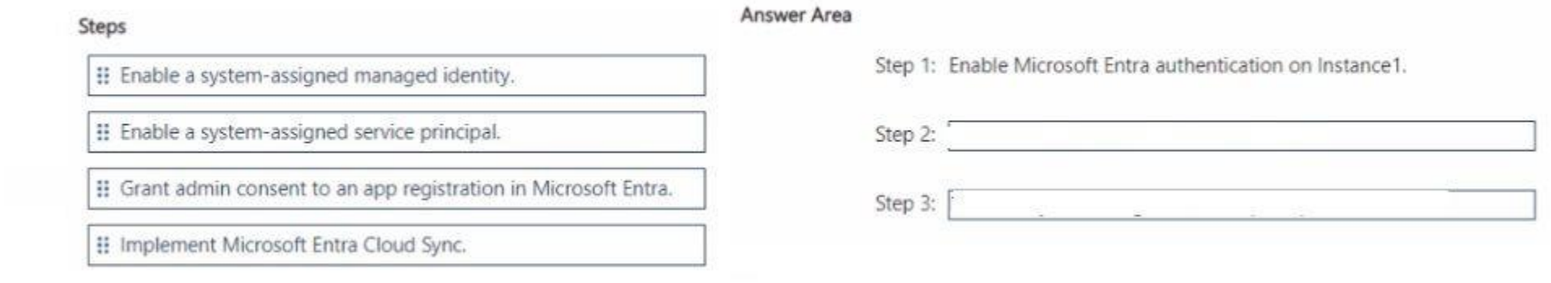

Correct Answer:

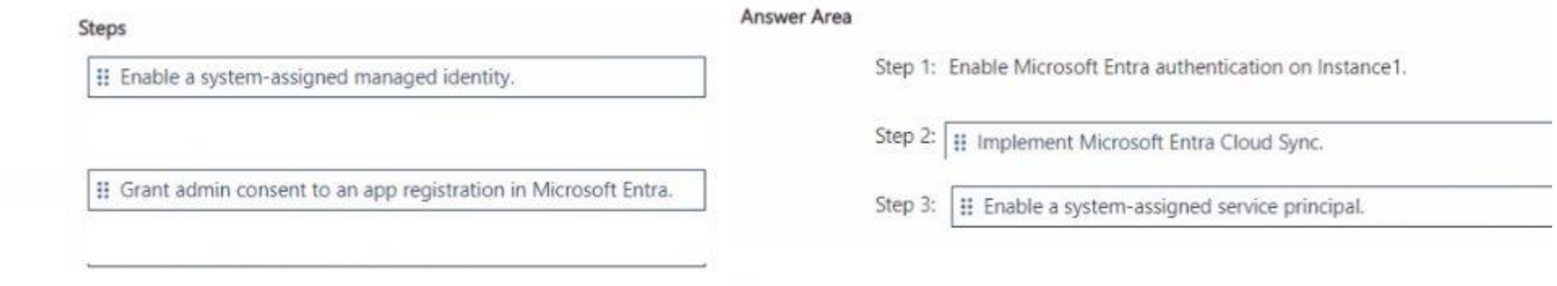

Section:
Explanation:

QUESTION 7
HOTSPOT
You need to recommend a service tier and a method to offload analytical workloads for the databases migrated from SVR1. The solution must meet the availability and business requirements.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

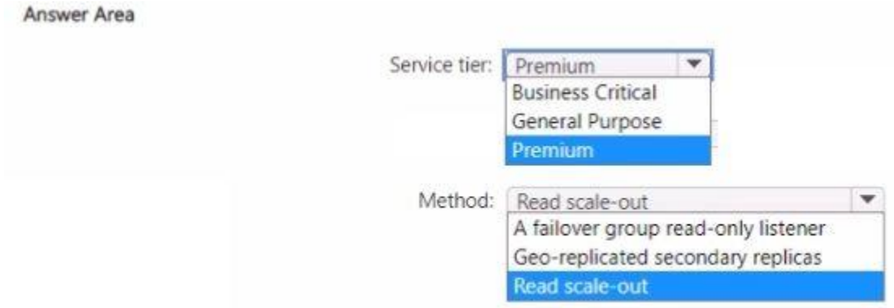

Answer Area:

Section:
Explanation:

QUESTION 8

You need to recommend a process to automate the management of DB3. The solution must meet the management requirements. What should be the first step of the process?

A. Configure Microsoft Entra authentication for the logical server that hosts DB3.
B. Create a database that has database-scoped credentials.
C. Configure a private endpoint for connectivity to DB3.
D. Create data base-scoped credentials in DB3.

Correct Answer: C
Section:

QUESTION 9
You need to identify the event_flle target for monitonng DB3 after the migration to Azure SQL Database. The solution must meet the management requirements,
What should you use as the event_file target?

A. an Azure SQL database
B. an Azure Blob Storage container
C. a SQL Server filegroup
D. an Azure Files share

Correct Answer: B
Section:

QUESTION 10
You need to identify the event_file target for monitonng DB3 after the migration to Azure SQL Database. The solution must meet the management requirements.
What should you use as the event_file target?

A. an Azure SQL database
B. an Azure Blob Storage container
C. a SQL Server filegroup
D. an Azure Files share

Correct Answer: B
Section:

QUESTION 11
HOTSPOT
You need to recommend which service and target endpoint to use when migrating the databases from SVR1 to Instance1. The solution must meet the availability requirements.
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:

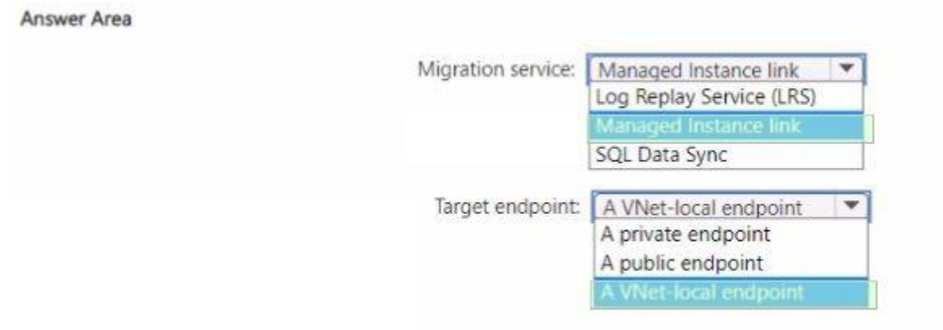

Section:
Explanation:

QUESTION 12
You need to recommend a solution that will enable remote developers to access DB1 and DB2. The solution must support the planned changes and meet the secunty requirements.
What should you include in the recommendation?

A. a public endpoint via a database-level firewall rule
B. a private endpoint
C. a public endpoint via a server-level firewall rule
D. a Point-to-Site (P2S) VPN

Correct Answer: B
Section:

QUESTION 13
You need to recommend a solution to ensure that the performance of DB3 is optimized after the migration to Azure SQL Database. The solution must meet availability requirements.
What should you include in the recommendation?

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

Correct Answer: C
Section:

QUESTION 14
You need to recommend a solution to meet the security requirements and the business requirements for DB3. What should you recommend as the first step of the solution?

A. Run the sys.5p_cdc_enble_db stored procedure.
B. Run the alter table statement and specify the enable chahgc_tracking clause.
C. Run the alter database statement and specify the set cmange_trackinc> - on clause.
D. Run the sp_addarticle stored procedure.

Correct Answer: C
Section:

04 - Monitor and Optimize Operational Resources

QUESTION 1
You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 contains a table named CustomerPII. You need to record whenever users query the CustomerPII table.
Which two options should you enable? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. server audit specification
B. SQL Server audit
C. database audit specification
D. a server principal

Correct Answer: A, C
Section:
Explanation:
An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):
A server policy applies to all existing and newly created databases on the server.
If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings. Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.
Note:
The Server Audit Specification object belongs to an audit.
A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auditing-overview

QUESTION 2
A company plans to use Apache Spark analytics to analyze intrusion detection data.
You need to recommend a solution to analyze network and system activity data for malicious activities and policy violations. The solution must minimize administrative efforts. What should you recommend?

A. Azure Data Lake Storage
B. Azure Databricks
C. Azure HDInsight
D. Azure Data Factory

Correct Answer: C
Section:
Explanation:
Azure HDInsight offers pre-made, monitoring dashboards in the form of solutions that can be used to monitor the workloads running on your clusters. There are solutions for Apache Spark, Hadoop, Apache Kafka, live long and process (LLAP), Apache HBase, and Apache Storm available in the Azure Marketplace.
Note: With Azure HDInsight you can set up Azure Monitor alerts that will trigger when the value of a metric or the results of a query meet certain conditions. You can condition on a query returning a record with a value that is greater than or less than a certain threshold, or even on the number of results returned by a query. For example, you could create an alert to send an email if a Spark job fails or if a Kafka disk usage becomes over 90 percent full.
Reference:
https://azure.microsoft.com/en-us/blog/monitoring-on-azure-hdinsight-part-4-workload-metrics-and-logs/

QUESTION 3
You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1. Several users execute adhoc queries to DW1 concurrently.
You regularly perform automated data loads to DW1.
You need to ensure that the automated data loads have enough memory available to complete quickly and successfully when the adhoc queries run. What should you do?

A. Assign a smaller resource class to the automated data load queries.
B. Create sampled statistics to every column in each table of DW1.
C. Assign a larger resource class to the automated data load queries.
D. Hash distribute the large fact tables in DW1 before performing the automated data loads.

Correct Answer: C
Section:
Explanation:
The performance capacity of a query is determined by the user's resource class.
Smaller resource classes reduce the maximum memory per query, but increase concurrency. Larger resource classes increase the maximum memory per query, but reduce concurrency.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workload-management

QUESTION 4

You are monitoring an Azure Stream Analytics job.

You discover that the Backlogged input Events metric is increasing slowly and is consistently non-zero. You need to ensure that the job can handle all the events.

What should you do?

A. Remove any named consumer groups from the connection and use $default.
B. Change the compatibility level of the Stream Analytics job.
C. Create an additional output stream for the existing input stream.
D. Increase the number of streaming units (SUs).

Correct Answer: D
Section:
Explanation:
Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring

QUESTION 5
You have an Azure Stream Analytics job.
You need to ensure that the job has enough streaming units provisioned.
You configure monitoring of the SU % Utilization metric.
Which two additional metrics should you monitor? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Late Input Events
B. Out of order Events
C. Backlogged Input Events
D. Watermark Delay
E. Function Events

Correct Answer: C, D
Section:
Explanation:
To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact. Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently nonzero, you should scale out your job, by increasing the SUs.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring

QUESTION 6
You have an Azure Databricks resource.
You need to log actions that relate to changes in compute for the Databricks resource.
Which Databricks services should you log?

A. clusters
B. jobs
C. DBFS
D. SSH
E. workspace

Correct Answer: E
Section:
Explanation:
Cloud Provider Infrastructure Logs.
Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace. Customers, especially in the regulated industries, also need records on activities like:
User access control to cloud data storage
Cloud Identity and Access Management roles
User access to cloud network and compute
Azure Databricks offers three distinct workloads on several VM Instances tailored for your data analytics workflow—the Jobs Compute and Jobs Light Compute workloads make it easy for data engineers to build and execute jobs, and the All-Purpose Compute workload makes it easy for data scientists to explore, visualize, manipulate, and share data and insights interactively.
Reference:
https://databricks.com/blog/2020/03/25/trust-but-verify-with-databricks.html

QUESTION 7
You have an Azure virtual machine based on a custom image named VM1.
VM1 hosts an instance of Microsoft SQL Server 2019 Standard.
You need to automate the maintenance of VM1 to meet the following requirements:
Automate the patching of SQL Server and Windows Server.
Automate full database backups and transaction log backups of the databases on VM1.
Minimize administrative effort.
What should you do first?

A. Enable a system-assigned managed identity for VM1
B. Register VM1 to the Microsoft.Sql resource provider
C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1
D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Correct Answer: D
Section:
Explanation:

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extension-automate-management

QUESTION 8
You have SQL Server on an Azure virtual machine.
You need to add a 4-TB volume that meets the following requirements:
Maximizes IOPs
Uses premium solid state drives (SSDs)
What should you do?

A. Attach two mirrored 4-TB SSDs.
B. Attach a stripe set that contains four 1-TB SSDs.
C. Attach a RAID-5 array that contains five 1-TB SSDs.
D. Attach a single 4-TB SSD.

Correct Answer: B
Section:
Explanation:

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

QUESTION 9
You have an Azure SQL database named db1 on a server named server1.
The Intelligent Insights diagnostics log identifies that several tables are missing indexes.
You need to ensure that indexes are created for the tables.
What should you do?

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

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview

QUESTION 10
You receive numerous alerts from Azure Monitor for an Azure SQL database.
You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period. Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Set Threshold Sensitivity to High
B. Set the Alert logic threshold to Dynamic
C. Set the Alert logic threshold to Static
D. Set Threshold Sensitivity to Low
E. Set Force Plan to On

Correct Answer: B, D
Section:
Explanation:
B: Dynamic Thresholds continuously learns the data of the metric series and tries to model it using a set of algorithms and methods. It detects patterns in the data such as seasonality (Hourly / Daily / Weekly), and is able to handle noisy metrics (such as machine CPU or memory) as well as metrics with low dispersion (such as availability and error rate). D: Alert threshold sensitivity is a high-level concept that controls the amount of deviation from metric behavior required to trigger an alert. Low - The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts. Incorrect Answers:
A: High - The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.
Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds

QUESTION 11
You have an Azure SQL database named sqldb1.
You need to minimize the amount of space by the data and log files of sqldb1.
What should you run?

A. DBCC SHRINKDATABASE
B. sp_clean_db_free_space

C. sp_clean_db_file_free_space
D. DBCC SHRINKFILE

Correct Answer: D
Section:
Explanation:
DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database.
Incorrect Answers:
D: To shrink one data or log file at a time for a specific database, execute the DBCC SHRINKFILE command.
Reference: https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql

QUESTION 12
You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.
The databases perform slower than expected.
You need to identify whether the performance issue relates to the use of tempdb by Azure SQL databases on sqlsrv1. What should you do?

A. Run Query Store-based queries
B. Review information provided by SQL Server Profiler-based traces
C. Review information provided by Query Performance Insight
D. Run dynamic management view-based queries

Correct Answer: C
Section:

QUESTION 13
You have an Azure SQL database named sqldb1.
You need to minimize the possibility of Query Store transitioning to a read-only state.
What should you do?

A. Halve the value of Data Flush Interval.
B. Double the value of Statistics Collection Interval.
C. Halve the value of Statistics Collection Interval
D. Double the value of Data Flush Interval.

Correct Answer: B
Section:
Explanation:

QUESTION 14
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory. You scale up the virtual machine to 16 vCPUSs and 64 GB of memory.
You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?

A. 2
B. 4
C. 8
D. 64

Correct Answer: D
Section:
Explanation:
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference: https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

QUESTION 15
You have 50 Azure SQL databases.
You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure. What should you do?

A. Create a diagnostic setting for the activity log that has the Security log enabled.
B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
C. Create an alert rule that uses a Metric signal type.
D. Create an alert rule that uses an Activity Log signal type.

Correct Answer: D
Section:
Explanation:
Activity log events - An alert can trigger on every event, or, only when a certain number of events occur. Incorrect Answers:
C: Metric values - The alert triggers when the value of a specified metric crosses a threshold you assign in either direction. That is, it triggers both when the condition is first met and then afterwards when that condition is no longer being met.
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/alerts-insights-configure-portal

QUESTION 16
You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1. You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server. Which resource type should you use to create the alert?

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

Correct Answer: C
Section:
Explanation:
There are resource types related to application code, compute infrastructure, networking, storage + databases. You can deploy up to 800 instances of a resource type in each resource group.
Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types

QUESTION 17
You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory. You scale up the virtual machine to 8 vCPUSs and 64 GB of memory.
You need to provide the lowest latency for tempdb.
What is the total number of data files that tempdb should contain?

A. 2
B. 4
C. 8
D. 64

Correct Answer: C
Section:
Explanation:
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
Reference: https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database

QUESTION 18
DRAG DROP
You are building an Azure virtual machine.
You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.
You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.
You need to configure storage for the virtual machine to support the database.
Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Select and Place:

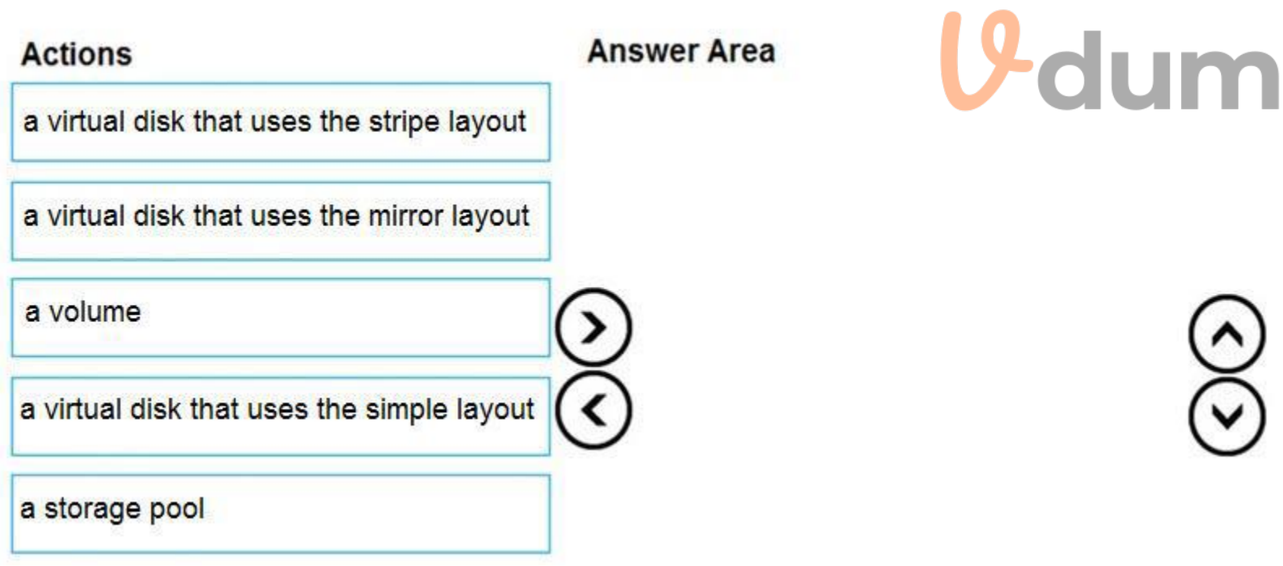

Correct Answer:

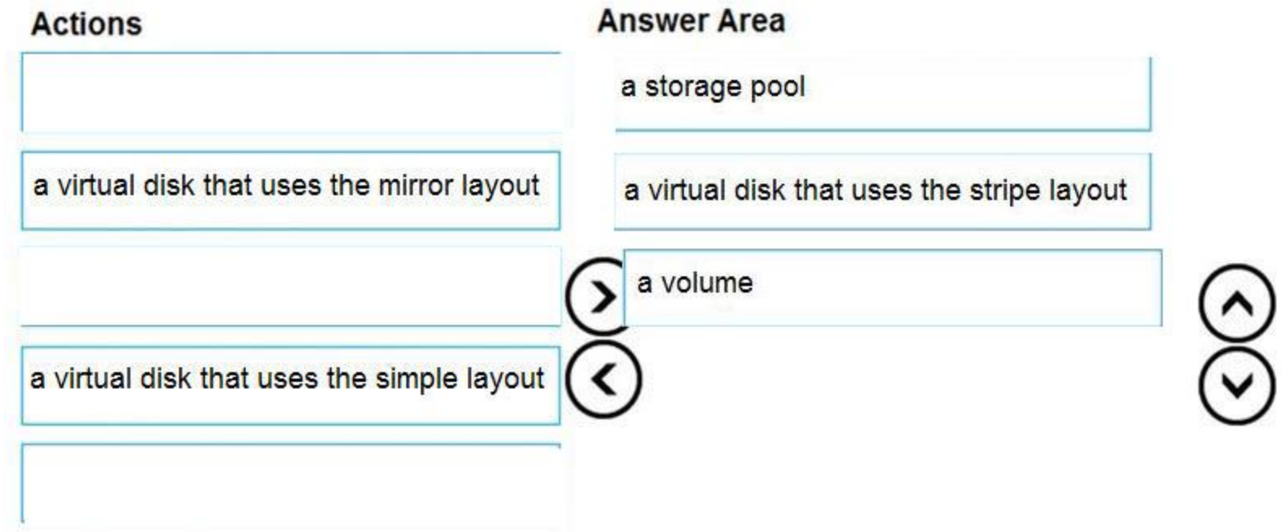

Section:
Explanation:
Follow these same steps to create striped virtual disk:
Create Log Storage Pool.
Create Virtual Disk
Create Volume
Box 1: a storage pool
Box 2: a virtual disk that uses stripe layout
Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.
Box 3: a volume
Reference:
https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/

QUESTION 19
HOTSPOT
You have an Azure SQL database named db1.
You need to retrieve the resource usage of db1 from the last week.
How should you complete the statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
SELECT *
FROM [                                  ▼]
     [sys.dm_db_resource_stats          ]
     [sys.dm_exec_requests              ]
     [sys.dm_user_db_resource_governance]
     [sys.resource_stats                ]
WHERE database_name = 'db1' AND
      start_time > [                    ▼] (day, -7, GETDATE())
                   [DATEADD             ]
                   [DATEDIFF            ]
                   [DATEPART            ]
                   [TODATETIMEOFFSET    ]
ORDER BY start_time DESC;
```


Answer Area:

Answer Area

```
SELECT *

FROM [                                  ▼]
     [sys.dm_db_resource_stats          ]
     [sys.dm_exec_requests              ]
     [sys.dm_user_db_resource_governance]
     [sys.resource_stats                ]

WHERE database_name = 'db1' AND

      start_time > [                    ▼] (day, -7, GETDATE())
                   [DATEADD             ]
                   [DATEDIFF            ]
                   [DATEPART            ]
                   [TODATETIMEOFFSET    ]

ORDER BY start_time DESC;
```


Section:
Explanation:
Box 1: sys.resource_stats
sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.
Box 2: DateAdd
The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

```
DECLARE @s datetime;
DECLARE @e datetime;
SET @s= DateAdd(d,-7,GetUTCDate());
SET @e= GETUTCDATE();
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization
FROM sys.resource_stats
WHERE start_time BETWEEN @s AND @e
GROUP BY database_name
HAVING AVG(avg_cpu_percent) >= 80
```

Incorrect Answers:
sys.dm_exec_requests:
sys.dm_exec_requests returns information about each request that is executing in SQL Server. It does not have a column named database_name.
sys.dm_db_resource_stats:

sys.dm_db_resource_stats does not have any start_time column.
Note: sys.dm_db_resource_stats returns CPU, I/O, and memory consumption for an Azure SQL Database database. One row exists for every 15 seconds, even if there is no activity in the database. Historical data is maintained for approximately one hour.
Sys.dm_user_db_resource_governance returns actual configuration and capacity settings used by resource governance mechanisms in the current database or elastic pool. It does not have any start_time column.
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-database

QUESTION 20
Your company uses Azure Stream Analytics to monitor devices.
The company plans to double the number of devices that are monitored.
You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load. Which metric should you monitor?

A. Input Deserialization Errors
B. Late Input Events
C. Early Input Events
D. Watermark delay

Correct Answer: D
Section:
Explanation:
The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.
The watermark delay metric can rise due to:
1. Not enough processing resources in Stream Analytics to handle the volume of input events.
2. Not enough throughput within the input event brokers, so they are throttled.
3. Output sinks are not provisioned with enough capacity, so they are throttled.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling

QUESTION 21
You manage an enterprise data warehouse in Azure Synapse Analytics.
Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries. You need to monitor resource utilization to determine the source of the performance issues.
Which metric should you monitor?

A. Local tempdb percentage
B. DWU percentage
C. Data Warehouse Units (DWU) used
D. Cache hit percentage

Correct Answer: A
Section:
Explanation:
Tempdb is used to hold intermediate results during query execution. High utilization of the tempdb database can lead to slow query performance. Note: If you have a query that is consuming a large amount of memory or have received an error message related to allocation of tempdb, it could be due to a very large CREATE TABLE AS SELECT (CTAS) or INSERT SELECT statement running that is failing in the final data movement operation.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-manage-monitor#monitor-tempdb

QUESTION 22
You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table. You need to identify the extent of the data skew in Table1.

What should you do in Synapse Studio?

A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
B. Connect to the built-in pool and run DBCC CHECKALLOC.
C. Connect to Pool1 and run DBCC CHECKALLOC.
D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Correct Answer: D
Section:
Explanation:
Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet

QUESTION 23
You have an Azure Synapse Analytics dedicated SQL pool.
You run PDW_SHOWSPACEUSED('dbo.FactInternetSales'); and get the results shown in the following table.

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

Which statement accurately describes the dbo.FactInternetSales table?

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

Correct Answer: D

Section:
Explanation:
The rows per distribution can vary up to 10% without a noticeable impact on performance. Here the distribution varies more than 10%. It is skewed.
Note: SHOWSPACEUSED displays the number of rows, disk space reserved, and disk space used for a specific table, or for all tables in a Azure Synapse Analytics or Parallel Data Warehouse database. This is a very quick and simple way to see the number of table rows that are stored in each of the 60 distributions of your database. Remember that for the most balanced performance, the rows in your distributed table should be spread evenly across all the distributions.
ROUND_ROBIN distributed tables should not be skewed. Data is distributed evenly across the nodes by design.
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute
https://github.com/rgl/azure-content/blob/master/articles/sql-data-warehouse/sql-data-warehouse-manage-distributed-data-skew.md

QUESTION 24
HOTSPOT
You are building an Azure Stream Analytics job to retrieve game data.
You need to ensure that the job returns the highest scoring record for each five-minute time interval of each game.
How should you complete the Stream Analytics query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

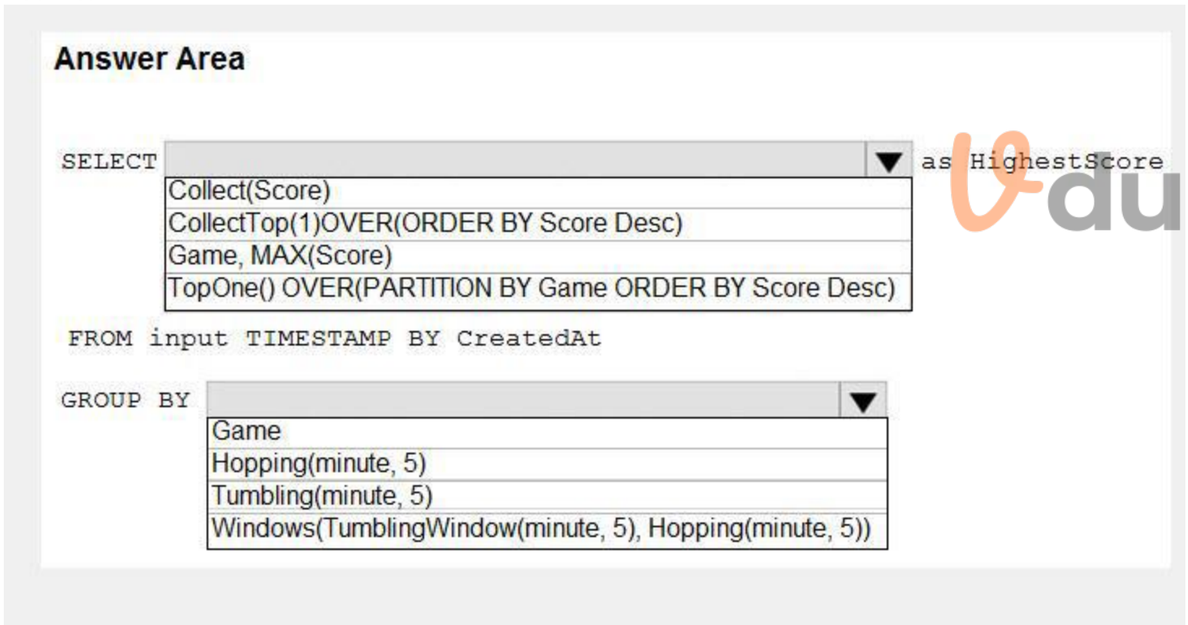

Answer Area:

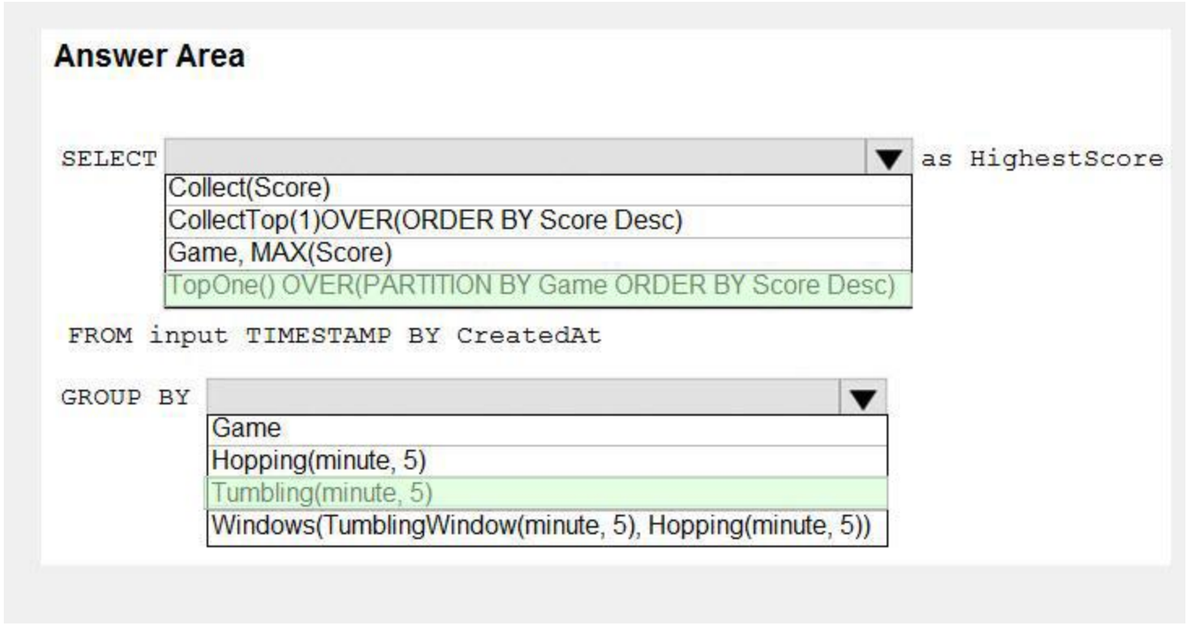

Section:

Explanation:

Box 1: TopOne() OVER(PARTITION BY Game ORDER BY Score Desc)

TopOne returns the top-rank record, where rank defines the ranking position of the event in the window according to the specified ordering. Ordering/ranking is based on event columns and can be specified in ORDER BY clause.

Analytic Function Syntax:

TopOne() OVER ([<PARTITION BY clause>] ORDER BY (<column name> [ASC |DESC])+ <LIMIT DURATION clause> [<WHEN clause>])

Box 2: Tumbling(minute 5)

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

Tell me the count of Tweets per time zone every 10 seconds

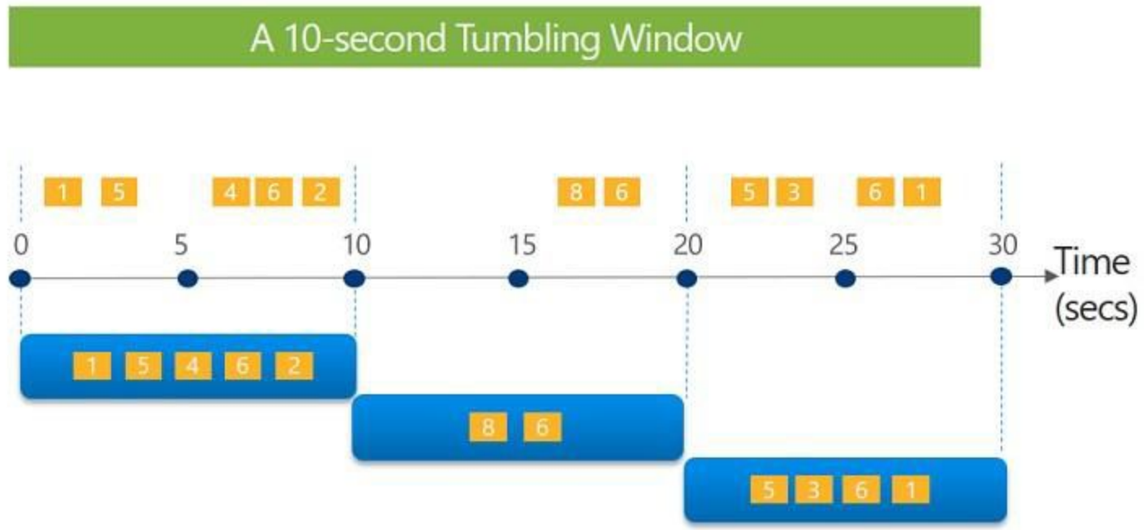


```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Reference:
https://docs.microsoft.com/en-us/stream-analytics-query/topone-azure-stream-analytics
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-functions.md

QUESTION 25
DRAG DROP
Your company analyzes images from security cameras and sends alerts to security teams that respond to unusual activity. The solution uses Azure Databricks.
You need to send Apache Spark level events, Spark Structured Streaming metrics, and application metrics to Azure Monitor.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions in the answer area and arrange them in the correct order.

Select and Place:

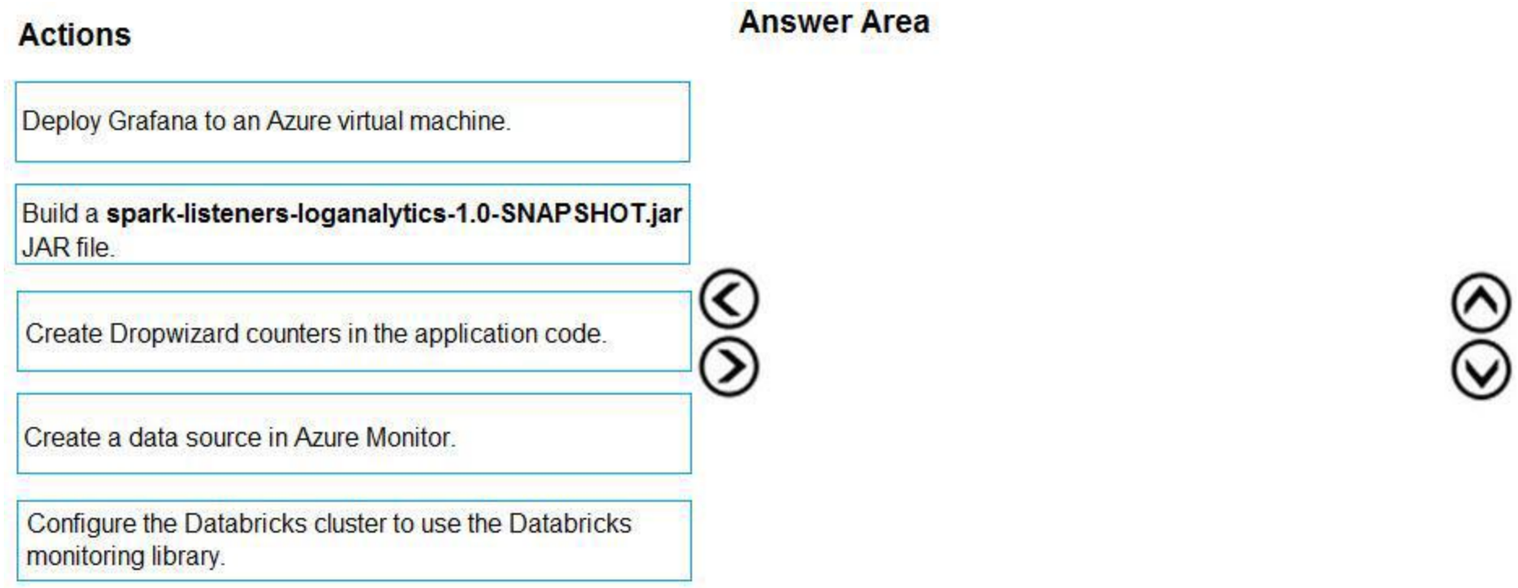

Correct Answer:

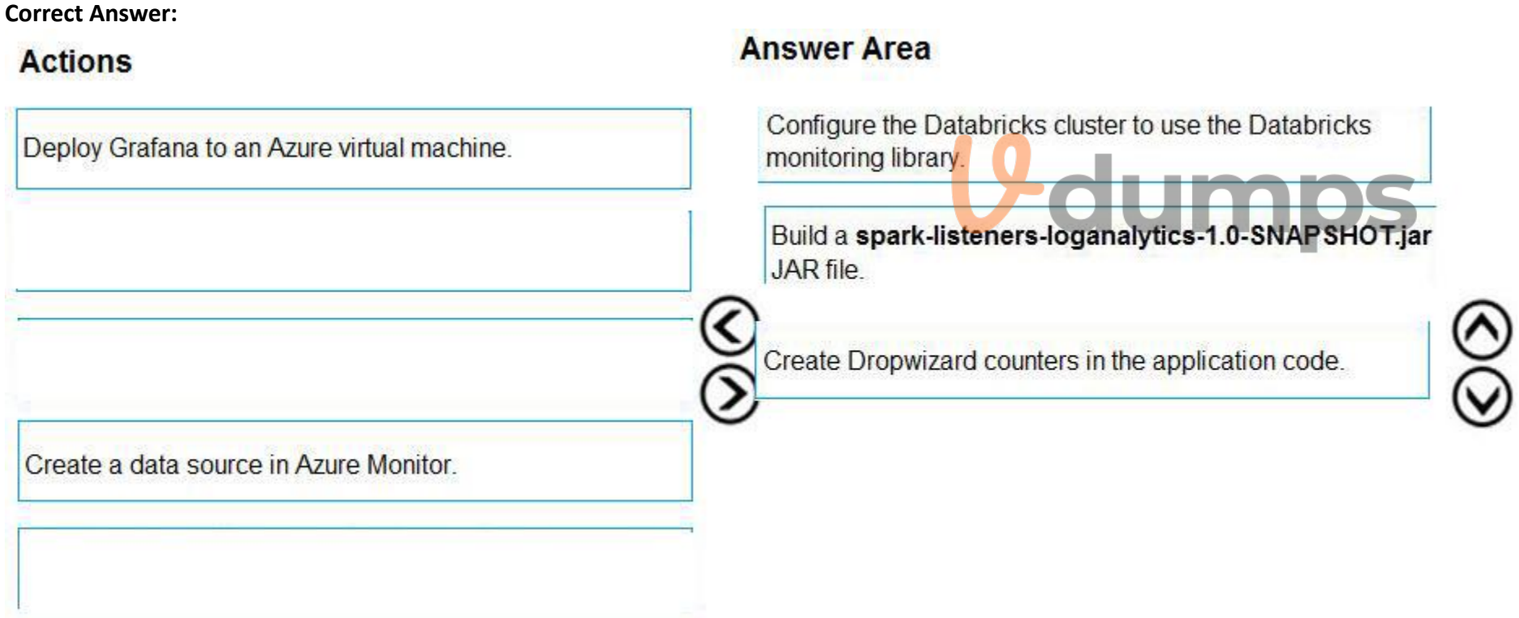

Section:
Explanation:
Send application metrics using Dropwizard.
Spark uses a configurable metrics system based on the Dropwizard Metrics Library.
To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps:
Step 1: Configure your Azure Databricks cluster to use the Databricksmonitoring library.
Prerequisite: Configure your Azure Databricks cluster to use the monitoring library.
Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file
Step 3: Create Dropwizard counters in your application code
Create Dropwizard gauges or counters in your application code

QUESTION 26
DRAG DROP
You have an Azure SQL managed instance named SQLMI1 that has Resource Governor enabled and is used by two apps named App1 and App2.
You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

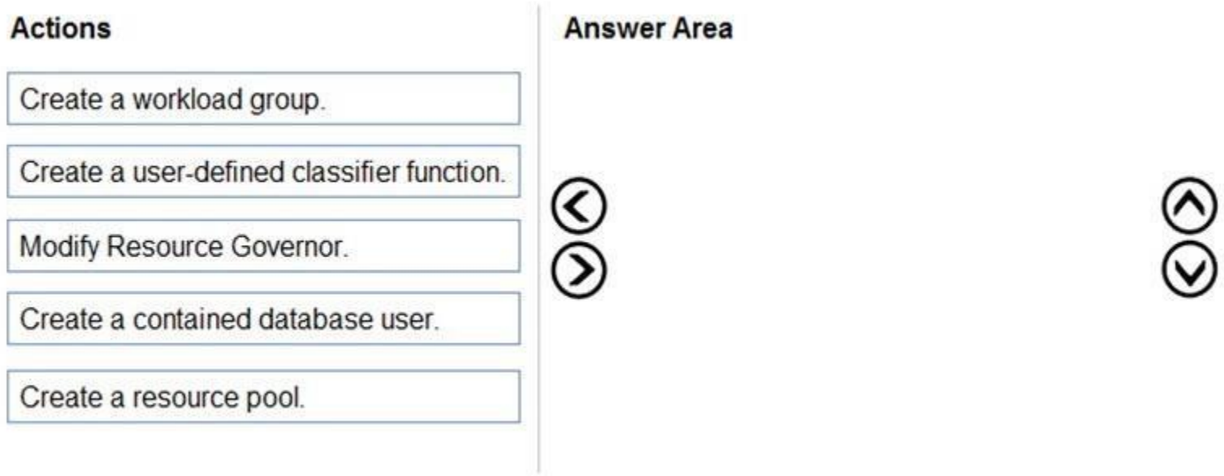

Correct Answer:

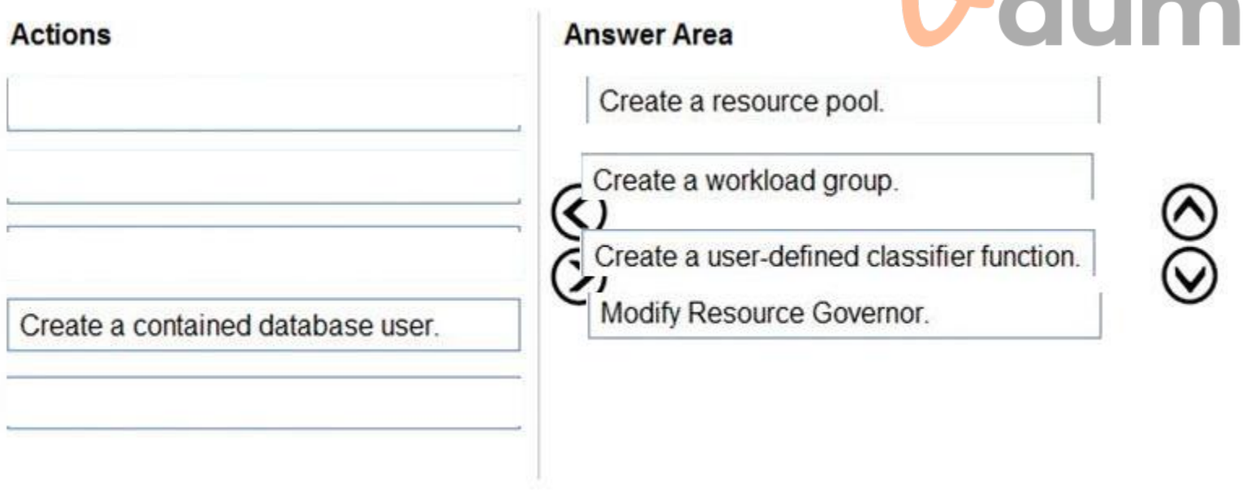

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15
https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-defined-function?view=sql-server-ver15

QUESTION 27
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance.
Solution: You shrink the transaction log file.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

QUESTION 28
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance.
Solution: You change the data file for the master database to autogrow by 10 percent.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

QUESTION 29
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance.
Solution: You reduce the use of table variables and temporary tables.
Does this meet the goal?

A. Yes

B. No

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

QUESTION 30
You have an Azure SQL database named db1 on a server named server1.
You need to modify the MAXDOP settings for db1.
What should you do?

A. Connect to db1 and run the sp_configure command.
B. Connect to the master database of server1 and run the sp_configure command.
C. Configure the extended properties of db1.
D. Modify the database scoped configuration of db1.

Correct Answer: D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism

QUESTION 31
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have SQL Server 2019 on an Azure virtual machine.
You are troubleshooting performance issues for a query in a SQL Server instance.
To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.
You need to improve system performance.
Solution: You create additional tempdb files.
Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention

QUESTION 32
You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.
You need to limit the number of IOPs that App2 queries generate on SQL1.
Which two actions should you perform on SQL1? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Correct Answer: B, D
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15

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.

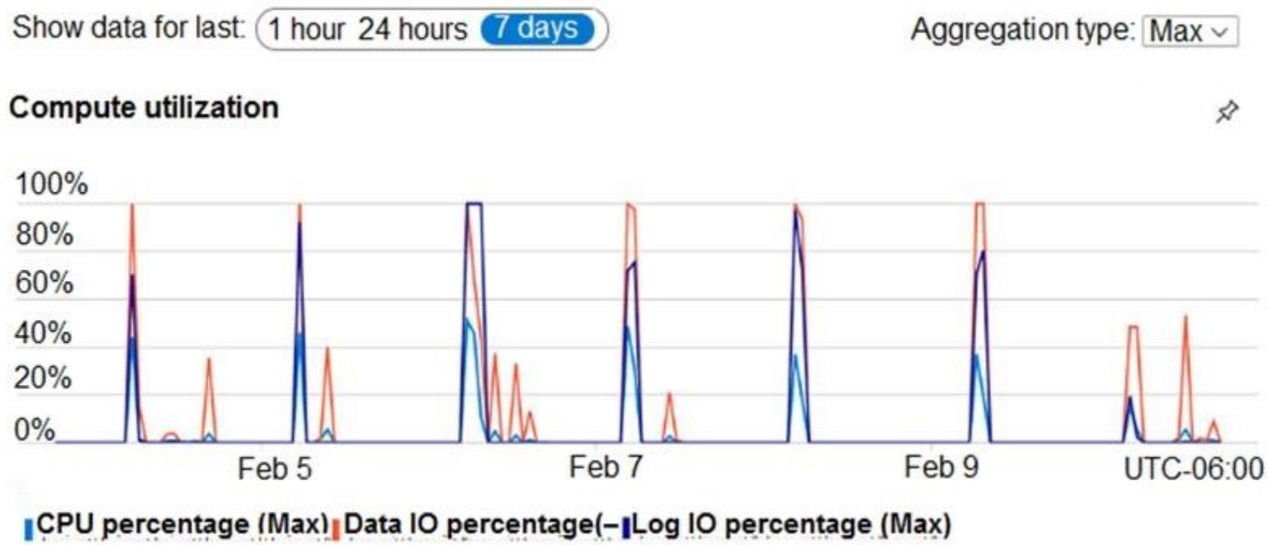

Role Assignments
For each customer's Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

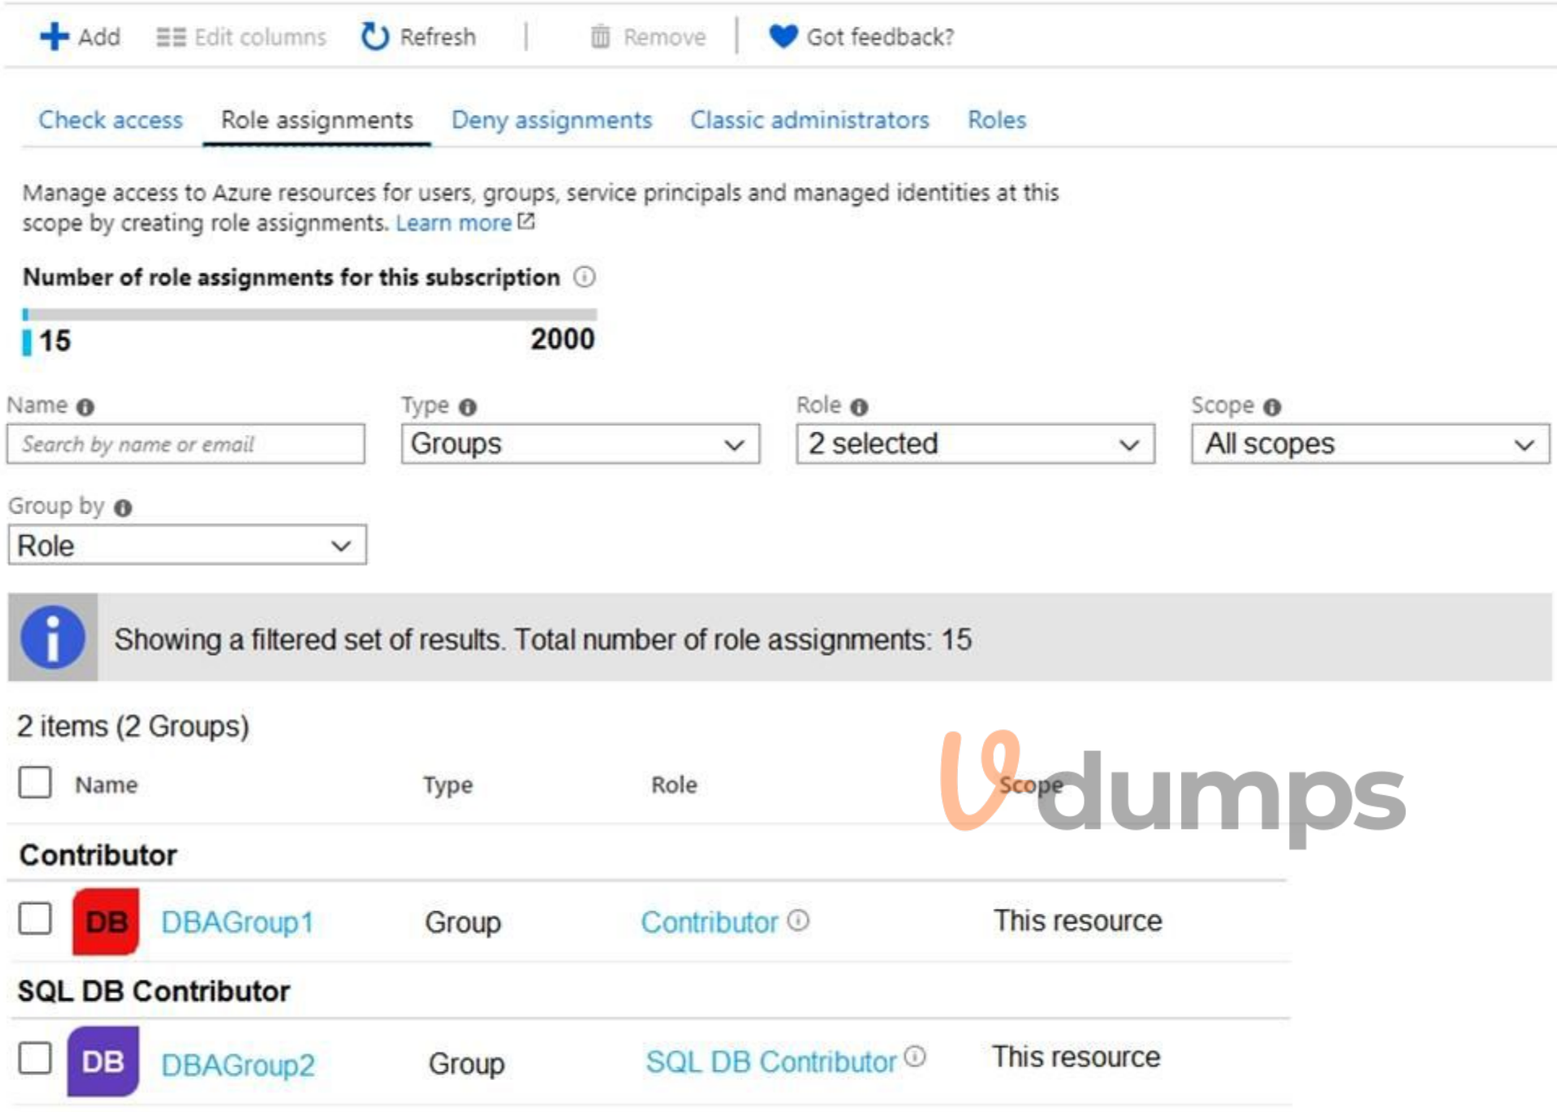

QUESTION 1
You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.
What should you do?

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

Create an Azure Automation account.

Import the SqlServer module.

Create a runbook that runs a PowerShell script.

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"

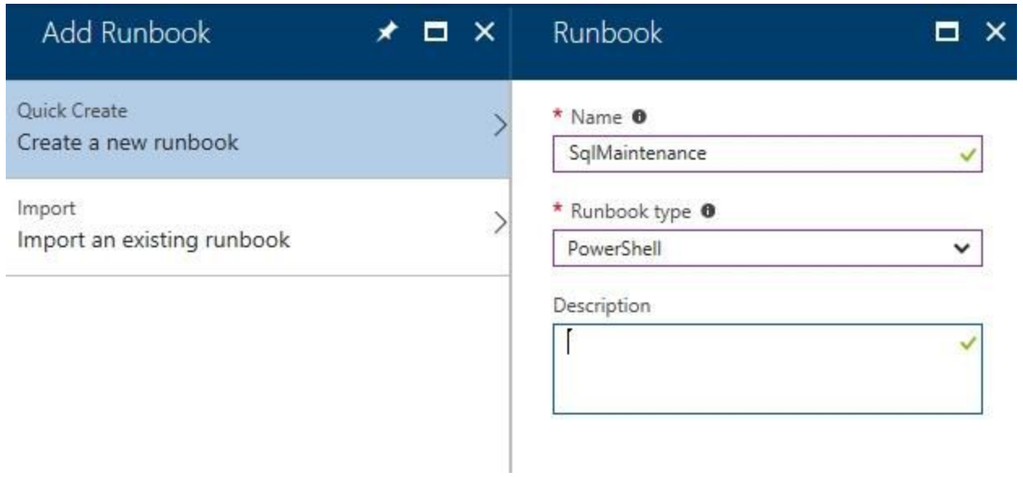

5. Schedule task (Step 4)
Steps:
1. Click on Schedules
2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.
Reference:
https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statistics-maintenance-using/ba-p/368974

03 - Perform Automation of Tasks

QUESTION 1
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 2
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:

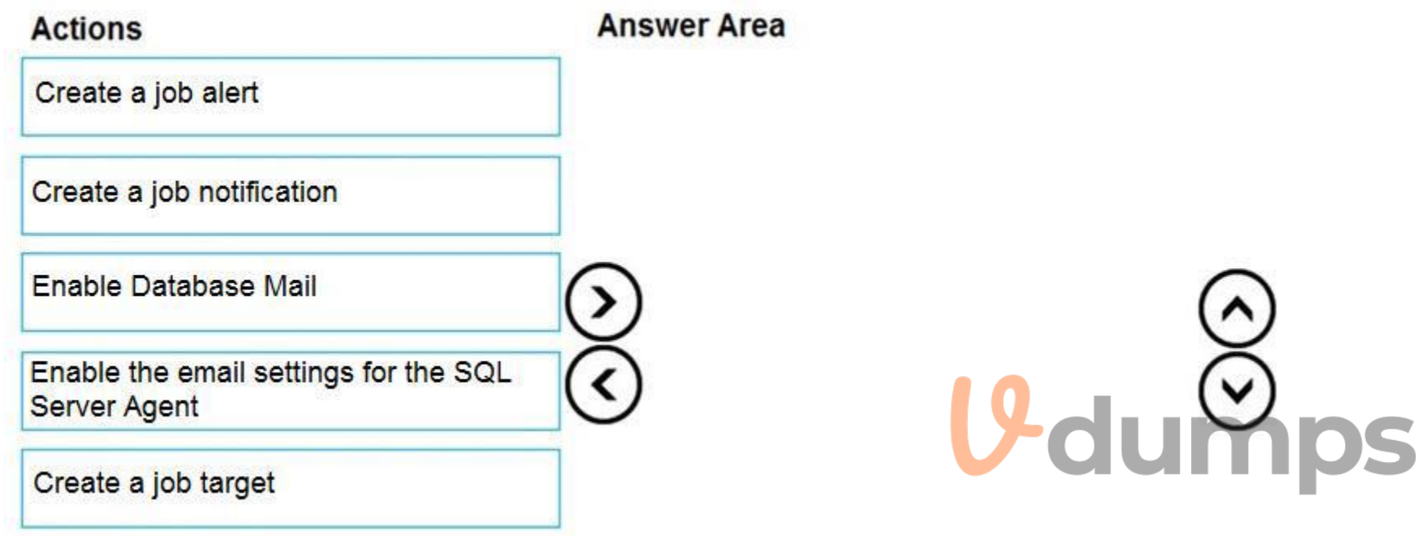

Correct Answer:

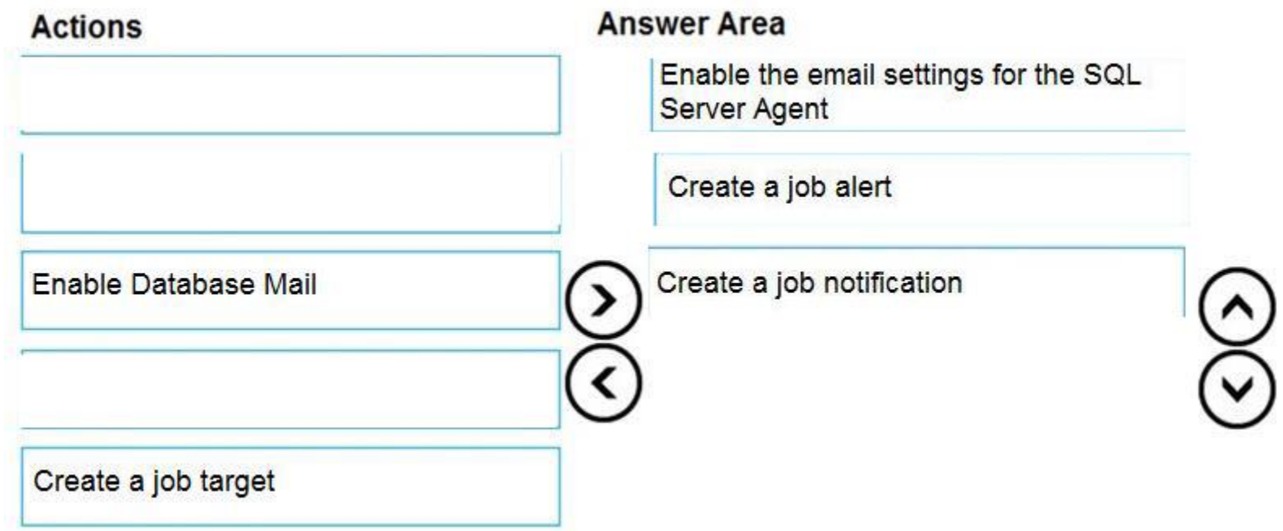

Section:
Explanation:

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 3
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:

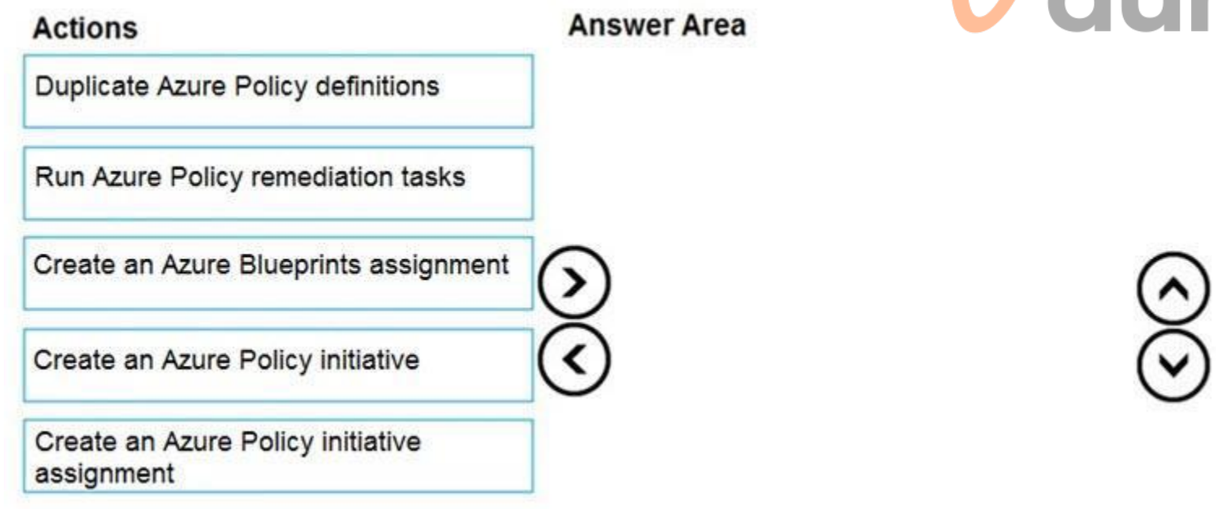

Correct Answer:

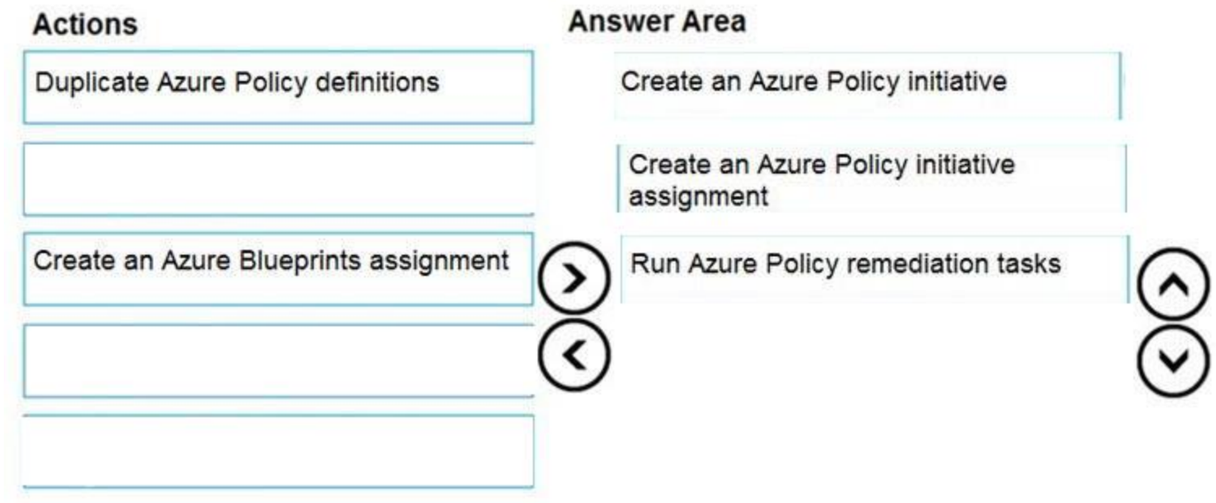

Section:
Explanation:
Step 1: Create an Azure Policy Initiative
The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.
With an initiative definition, you can group several 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 4
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 5
You have the following Azure Data Factory pipelines:
Ingest Data from System1
Ingest Data from System2
Populate Dimensions

Populate Facts
Ingest Data from System1 and Ingest Data from System2 have no dependencies. Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2. Populate Facts must execute after the Populate Dimensions pipeline. All the pipelines must execute every eight hours.
What should you do to schedule the pipelines for execution?

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

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

QUESTION 6
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

QUESTION 7
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 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 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 9
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 10
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 11
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 12
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:

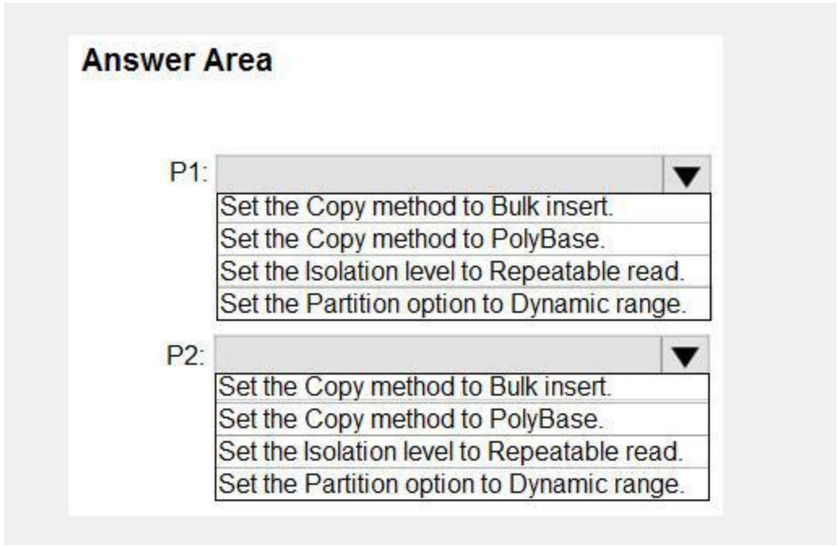

Answer Area:

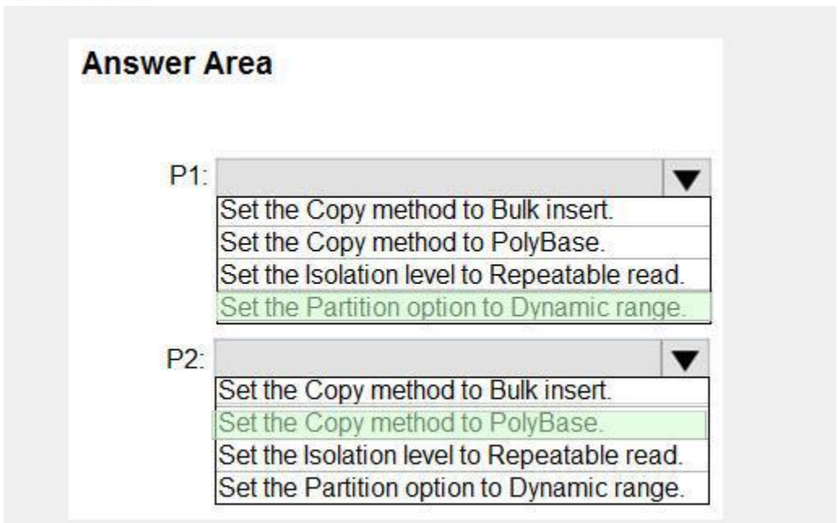

Section:
Explanation:
P1: Set the Partition option to Dynamic Range.
The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel.
P2: Set the Copy method to PolyBase
Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging 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 13

Hot Area:

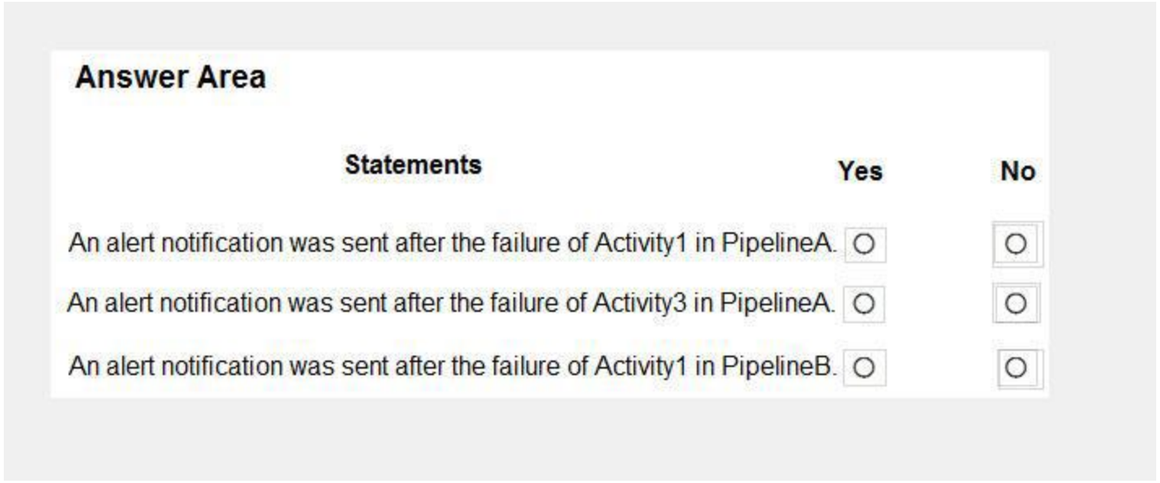

Answer Area:

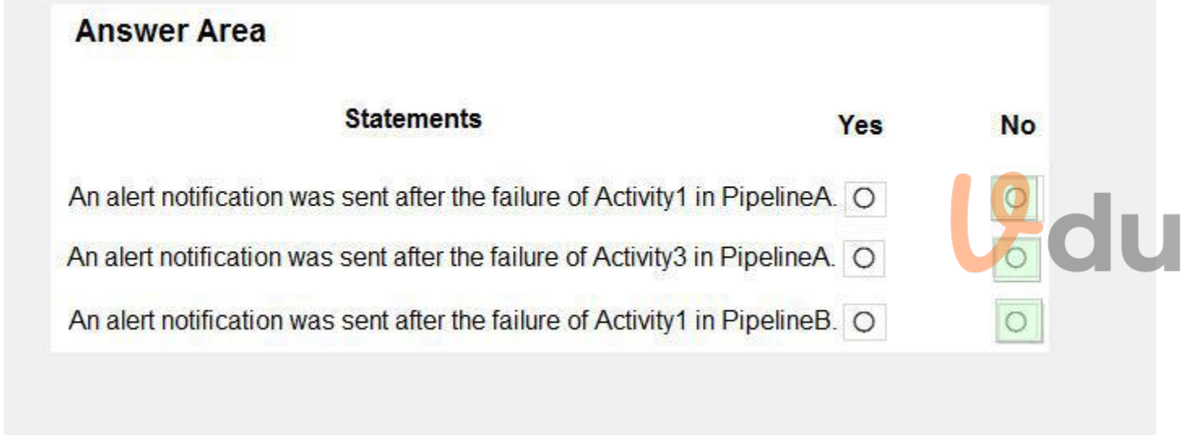

Section:
Explanation:
Box 1: No
Just one failure within the 5-minute interval.
Box 2: No
Just two failures within the 5-minute interval.
Box 3: No
Just two failures within the 5-minute interval.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview

QUESTION 14
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:

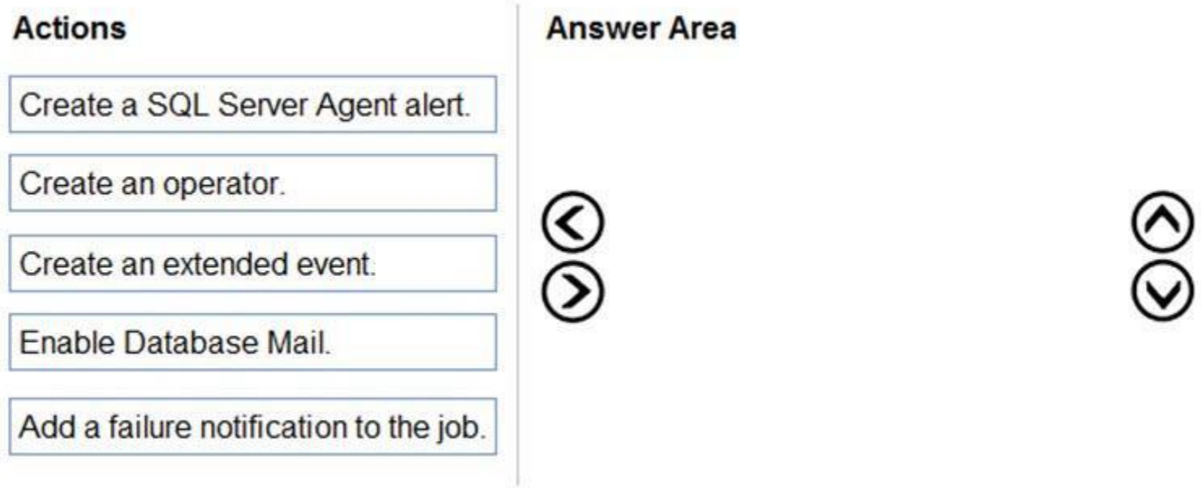

Correct Answer:

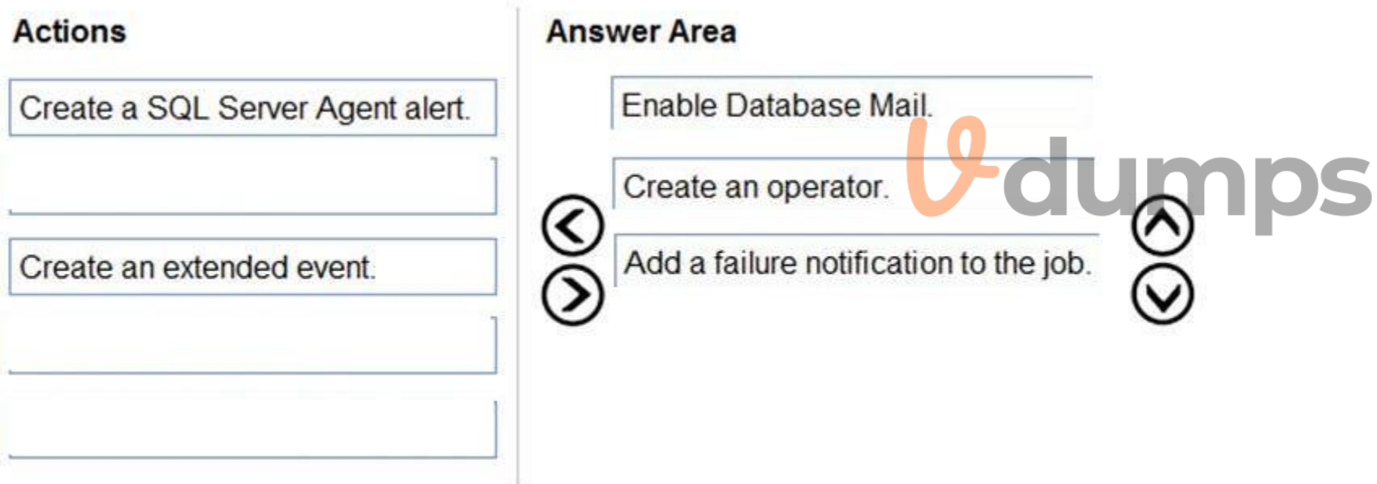

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/job-automation-managed-instance

QUESTION 15
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:

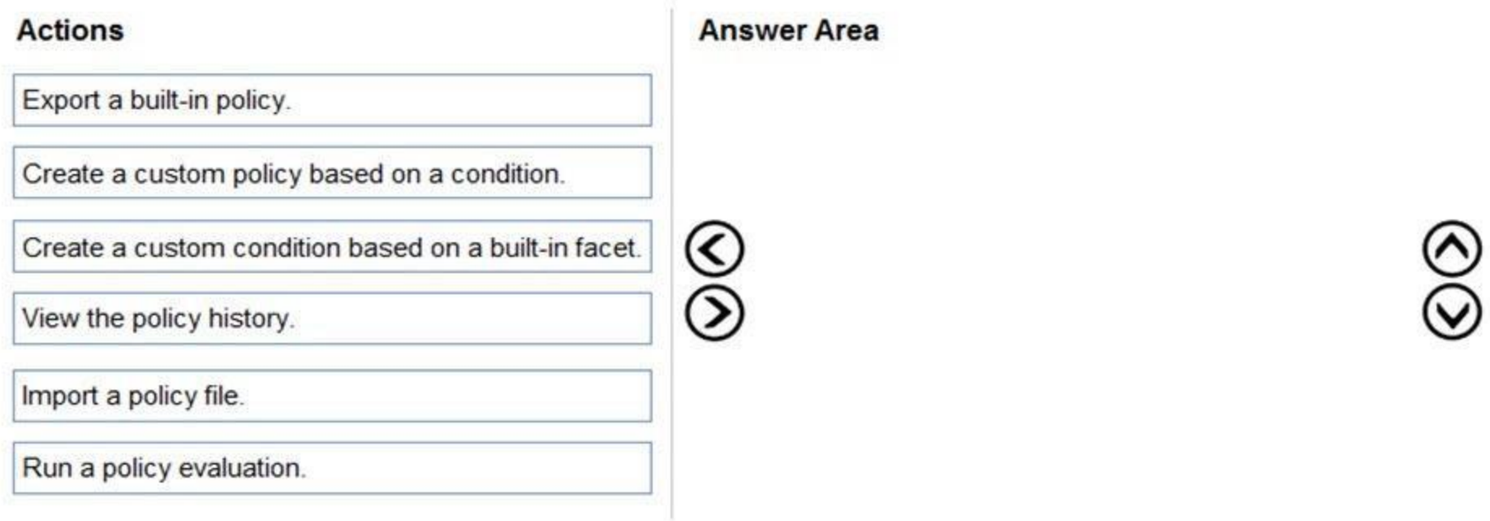

Correct Answer:

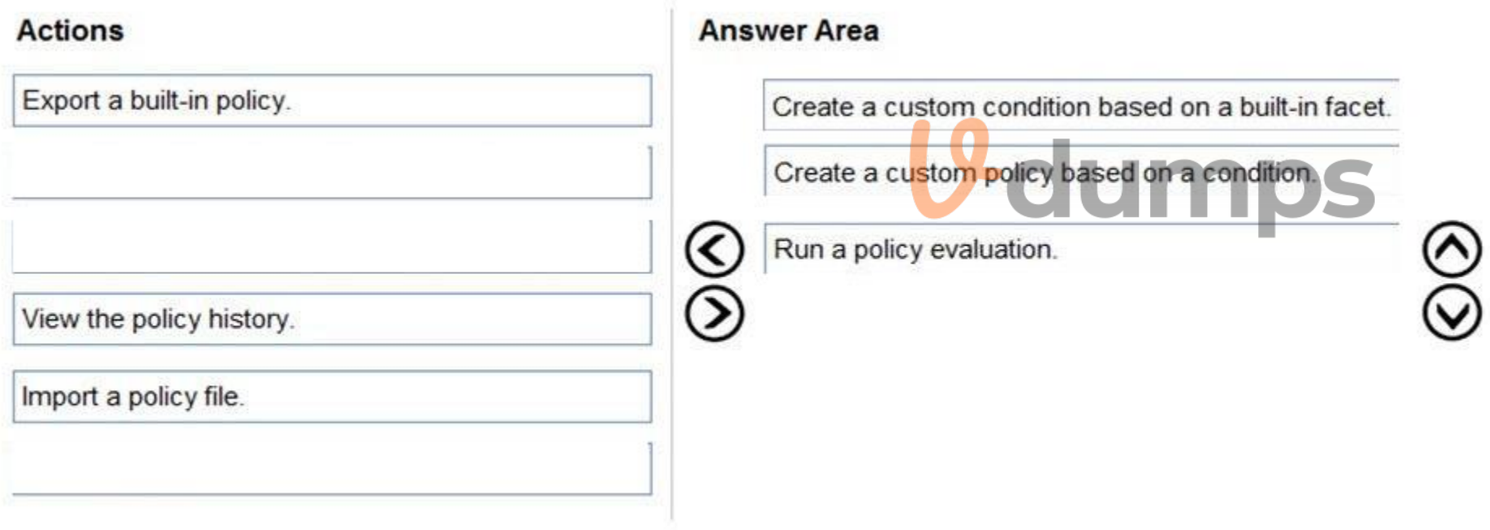

Section:
Explanation:
Reference:
https://www.mssqltips.com/sqlservertip/2298/enforce-sql-server-database-naming-conventions-using-policy-based-management/

QUESTION 16
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 17
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

Exam K

QUESTION 1
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:

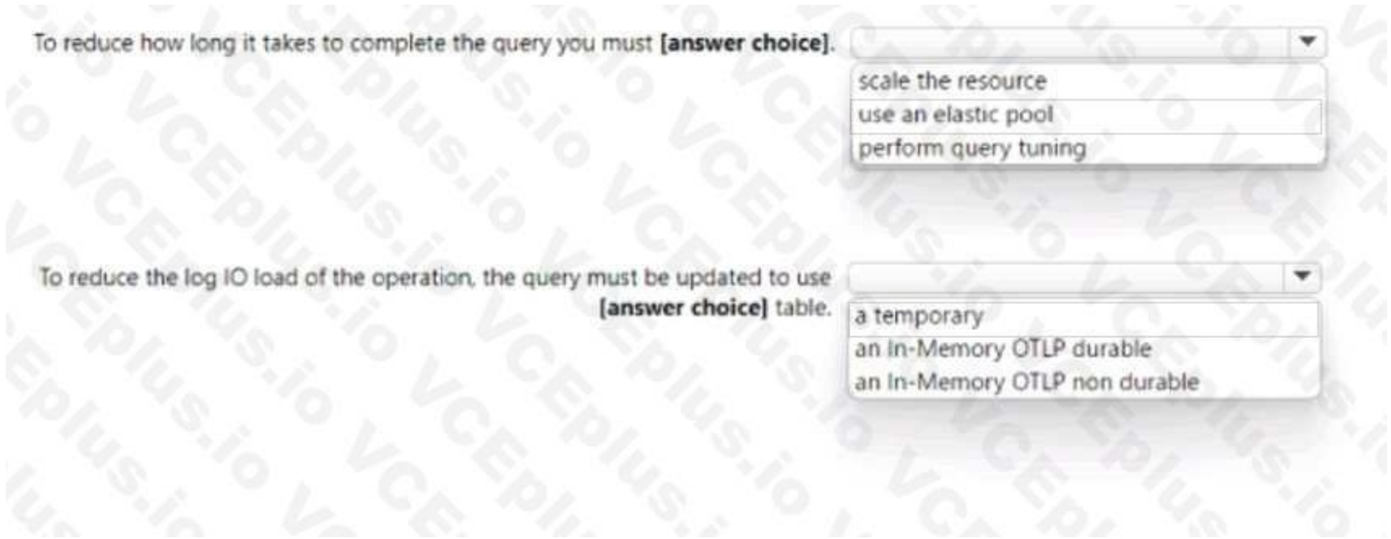

Answer Area:

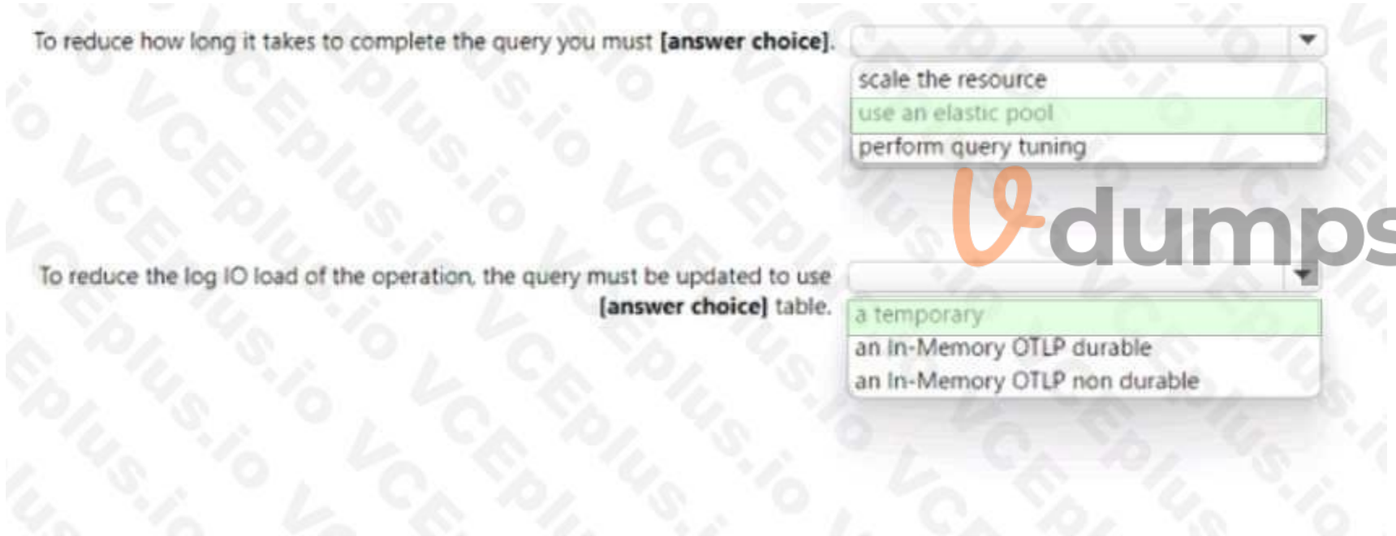

Section:
Explanation:

QUESTION 2
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:

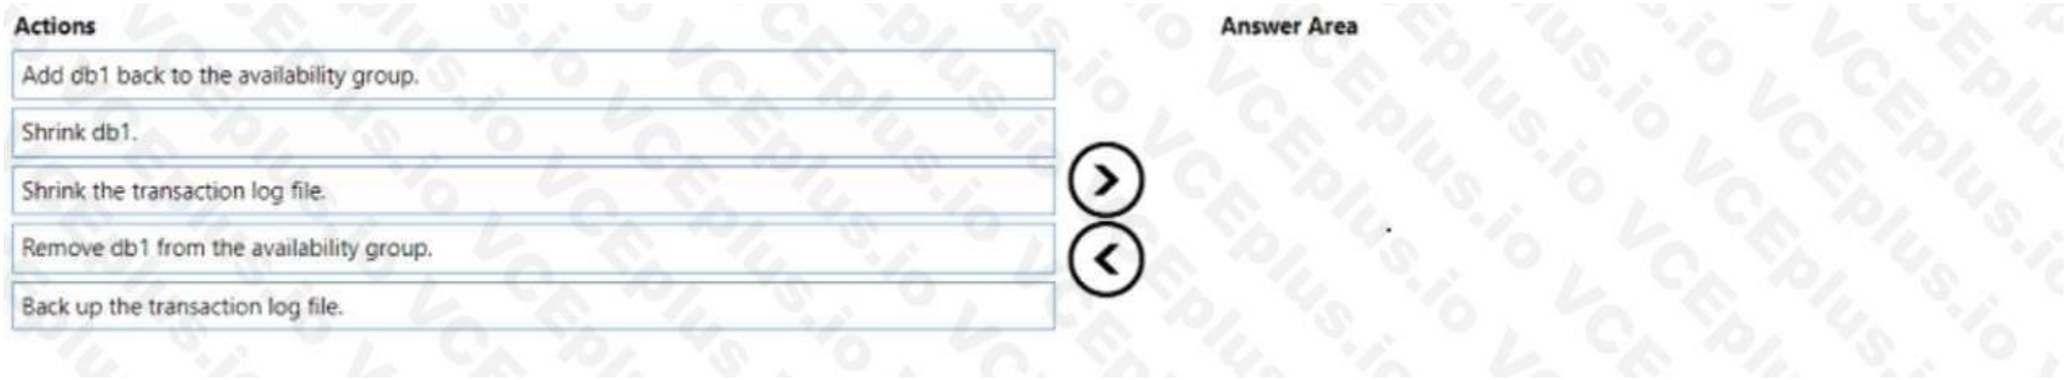

Correct Answer:

Section:
Explanation:

QUESTION 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:

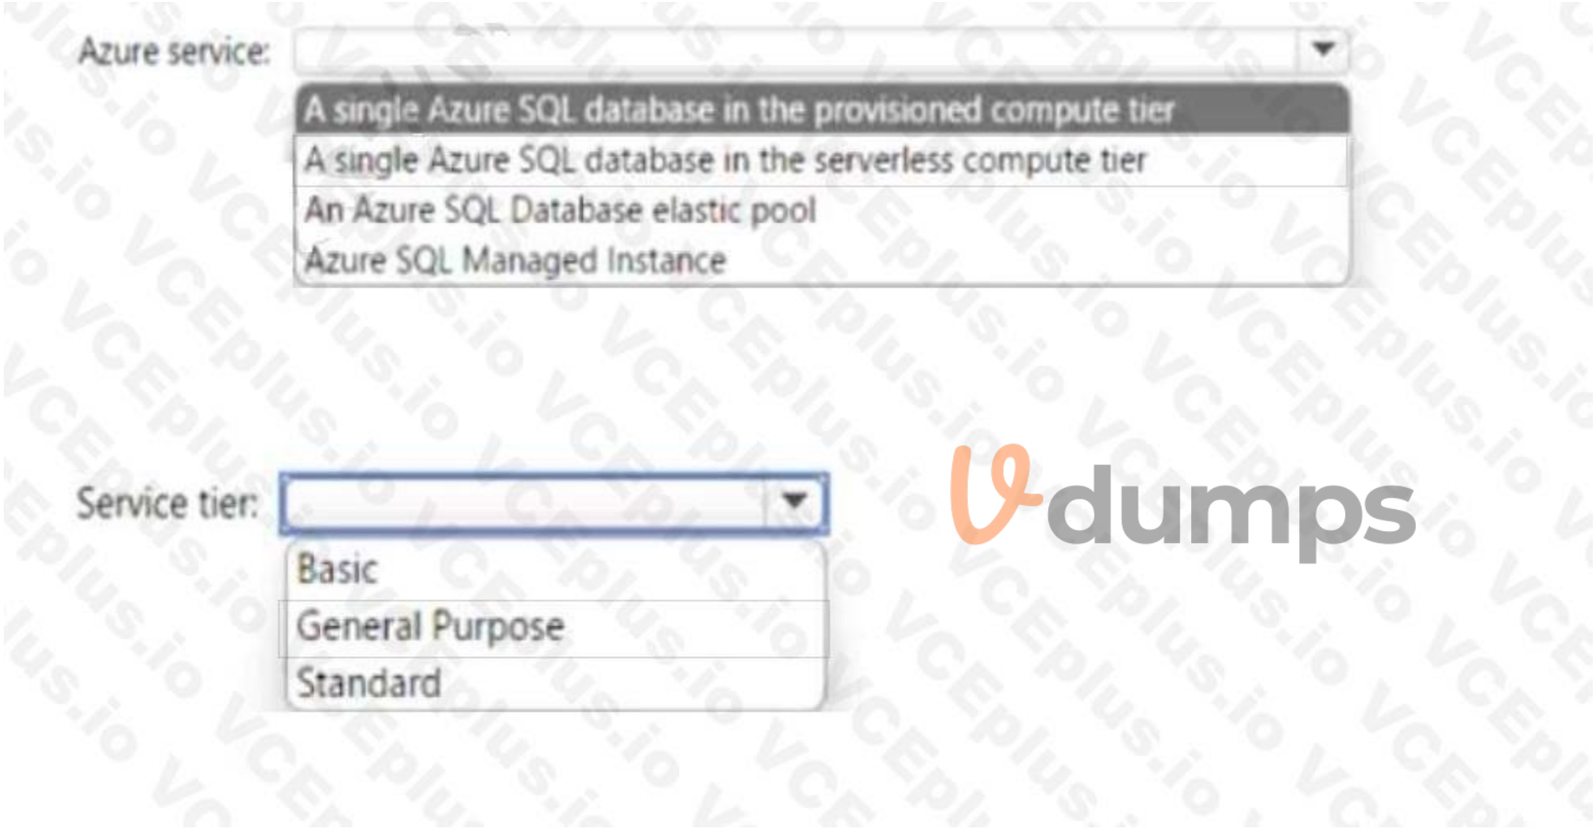

Answer Area:

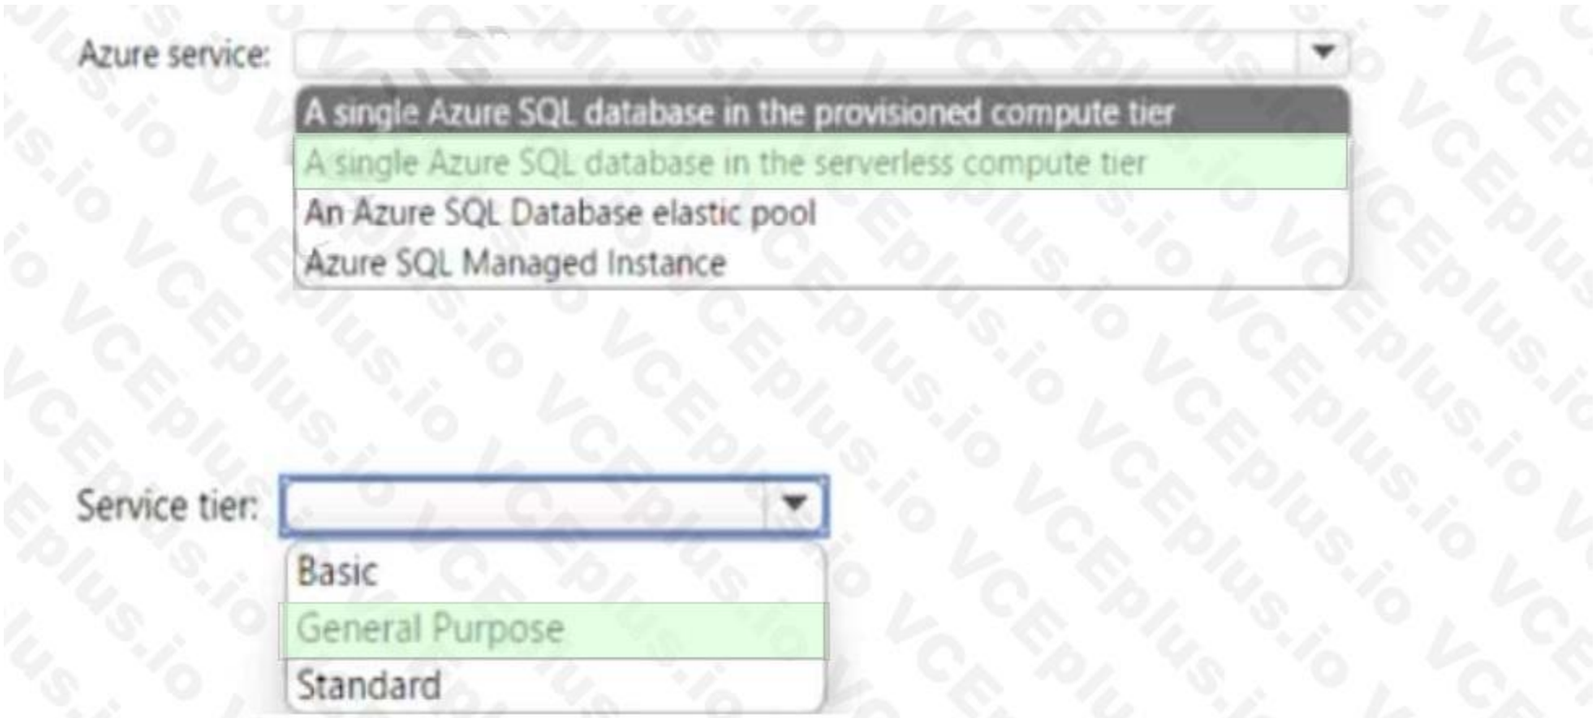

Section:
Explanation:

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
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 7
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 8
You have an Azure SQL database named DB1 that contains a private certificate named Sales. The private key for Sales is encrypted with a password. You need to change the password for the private key. Which Transact-SQL statement should you run?

A.
```
ALTER CERTIFICATE Sales
    WITH PRIVATE KEY (DECRYPTION BY PASSWORD = 'Mb^6BK&*w%',
    ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```
B.
```
ALTER CERTIFICATE Sales
    WITH PRIVATE KEY (ENCRYPTION BY PASSWORD = ' 6YY9YcD!pV');
```
C.
```
ALTER CERTIFICATE Sales       WITH PRIVATE KEY (FILE = 'D:\importkeys\SalesNew,      DECRYPTION BY PASSWORD = ' Mb^6BK&*w%');
```
D.
```
ALTER CERTIFICATE Sales       WITH PRIVATE KEY (DECRYPTION BY PASSWORD = ' EWYx9Xk+$#');
```

Correct Answer: C
Section:

QUESTION 9
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 10
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 11
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 12
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 13
You have an Azure virtual machine named Server1 that runs Windows Server 2022. Server! contains an instance of Microsoft SQL Server 2022 named SQL1 and a database named DB1.
You create a master key in the master database of SQL1.
You need to create an encrypted backup of DB1.
What should you do?

A. Create a symmetric key in DB1.
B. Enable visualization-based security (VBS) on Server1.
C. Create a certificate in DB1.

D. Create a certificate in the master database of SQLl.

Correct Answer: D
Section:

QUESTION 14
You have an Azure subscription that contains three instances of SQL Server on Azure Virtual Machines.
You plan to implement a disaster recovery solution.
You need to be able to perform disaster recovery drills regularly. The solution must meet the following requirements:
* Minimize administrative effort for the recovery drills.
* Isolate the recovery environment from the production environment.
What should you use?

A. Recovery Services vaults
B. Azure Site Recovery
C. Azure Backup
D. native Microsoft SQL Server backup

Correct Answer: B
Section:

QUESTION 15
You have two on-premises servers that run Windows Server 2019 and host a Microsoft SQL Server 2017 Always On availability group named AG1. AG1 contains a single database named DB1. You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Linux. You need to migrate DB1 to a SQL Server 2019 instance on VM1. The solution must minimize the downtime of DB1 during the migration. What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

To prepare for the migration: [ ]

To perform the migration, use: [ ]

A.

Answer Area

To prepare for the migration: Create a SQL Server 2019 Always On availability group on VM1.

To perform the migration, use: Azure Migrate

Correct Answer: A
Section:

QUESTION 16

DRAG DROP
You have an Azure subscription.
You plan to deploy a new Azure virtual machine that will host a Microsoft SQL Server instance.
You need to configure the disks on the virtual machine. The solution must meet the following requirements:
• Minimize latency for transaction logs.
• Minimize the impact on IO Of the virtual machine.
Which type of disk should you use for each workload? To answer. drag the appropriate disk types to the correct workloads. Each disk type may be used once, 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:

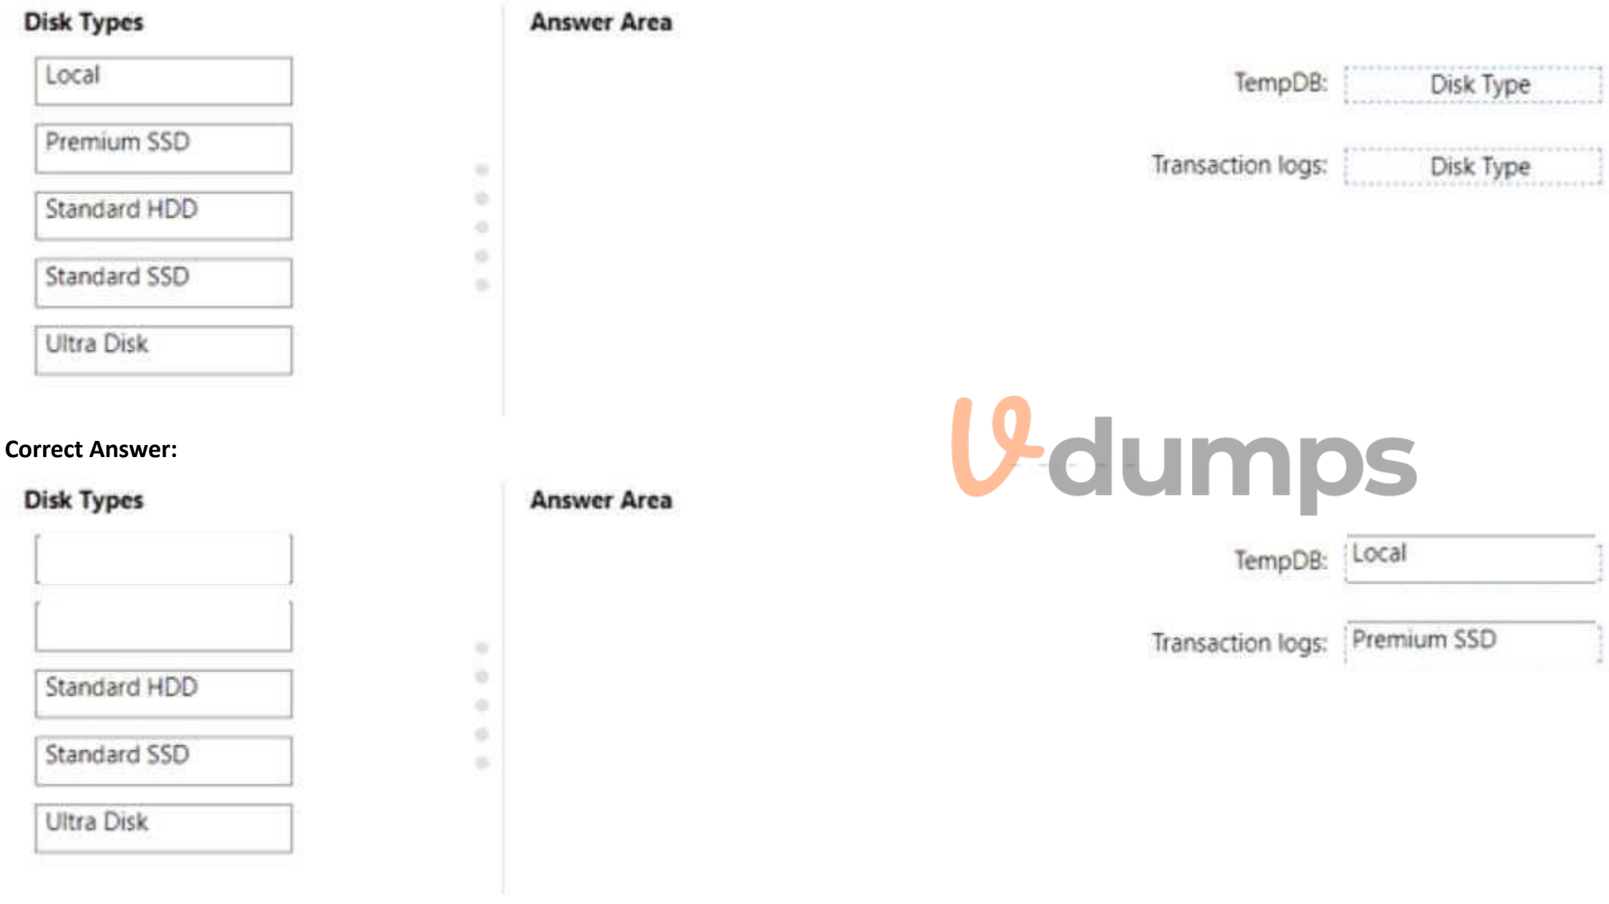

Correct Answer:

Section:
Explanation:

QUESTION 17
Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database named db1. You have an Azure subscription.
You plan to migrate db1 to an Azure SQL managed instance.
You need to create the SQL managed instance. The solution must minimize the disk latency of the instance. Which service tier should you use?

A. Hyperscale
B. General Purpose

C. Premium
D. Business Critical

Correct Answer: D
Section:

QUESTION 18
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:

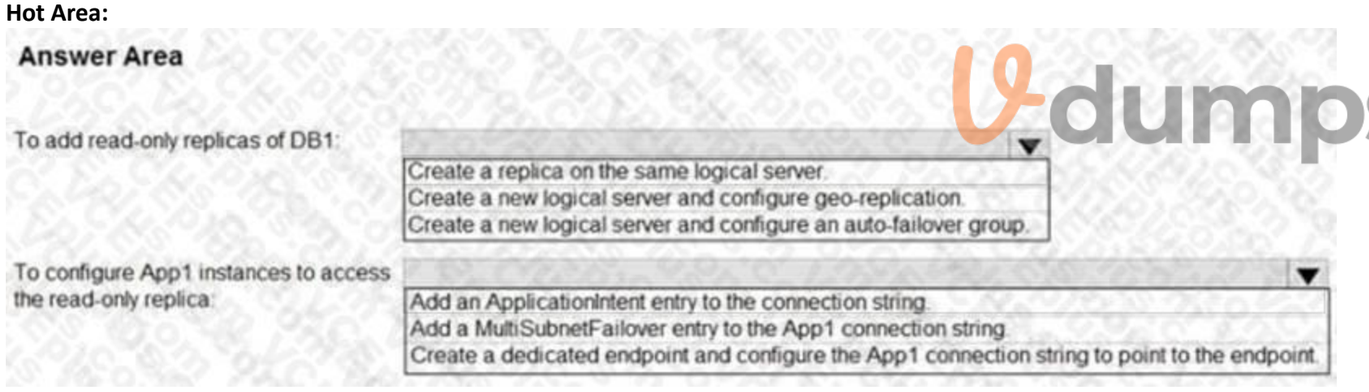

Answer Area:

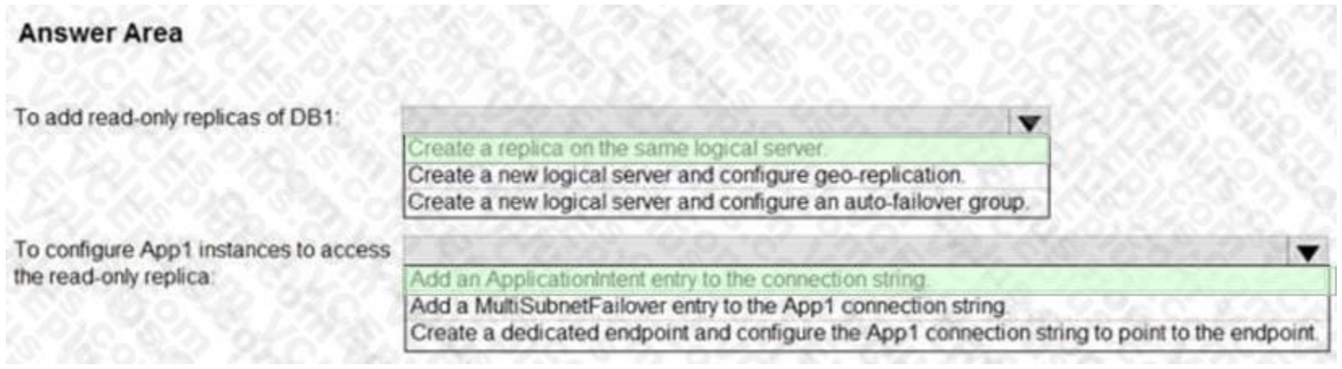

Section:
Explanation:
Reference:

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

QUESTION 19
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:

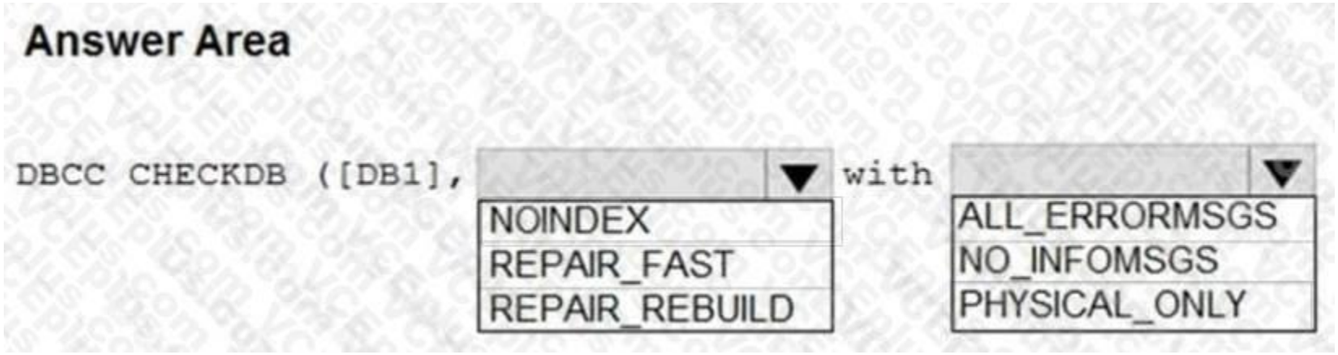

Answer Area:

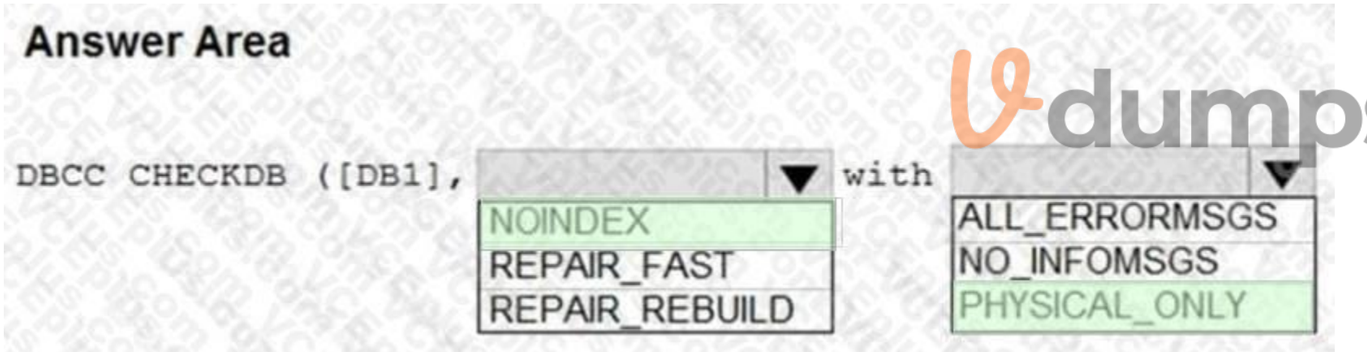

Section:
Explanation:
Reference: https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql?view=sql-server-ver15

QUESTION 20
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 21
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.

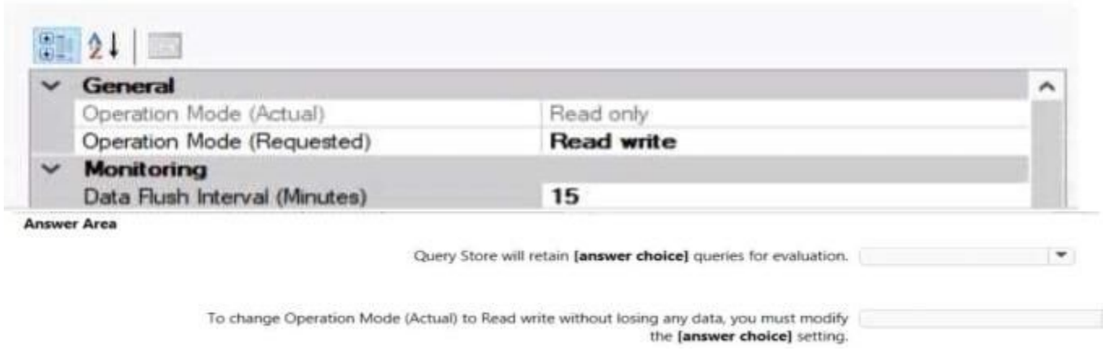

A.

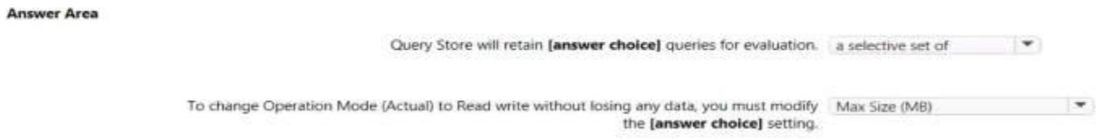

Correct Answer: A
Section:

QUESTION 22
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 23
You have an Azure subscription that contains a logical SQL server named Server1. The master database of Server1 contains a user named User1. You need to ensure that User1 can create databases on Server1. Which database role should you assign to User1?

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

Correct Answer: B
Section:

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:

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 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 App1 as soon as possible, without considering immediate costs. What should you do?

A. Increase the number Of vCores allocated to dbl.

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

Correct Answer: D
Section:

QUESTION 26
HOTSPOT
You have an Azure SQL database named DB 1 in the General Purpose service tier.
You need to monitor DB 1 by using SQL Insights.
What should you include in the solution? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

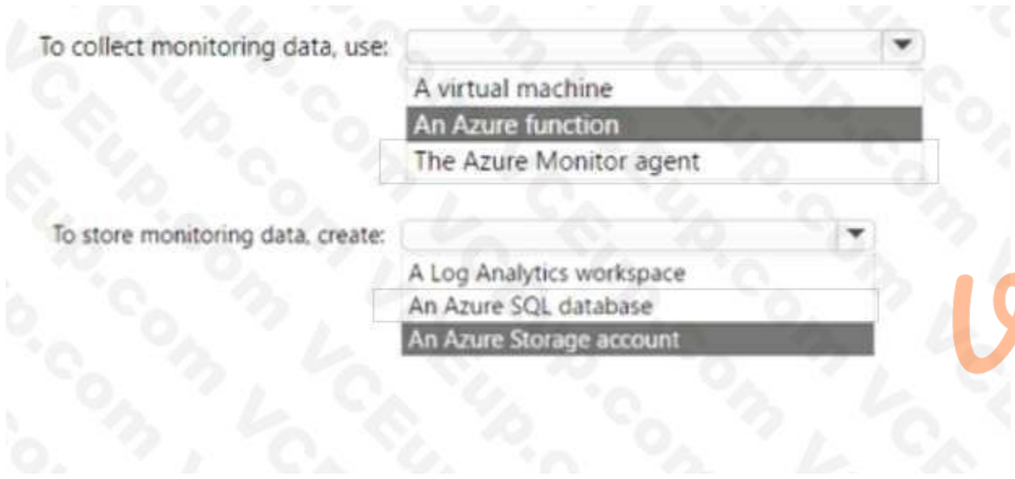

Answer Area:

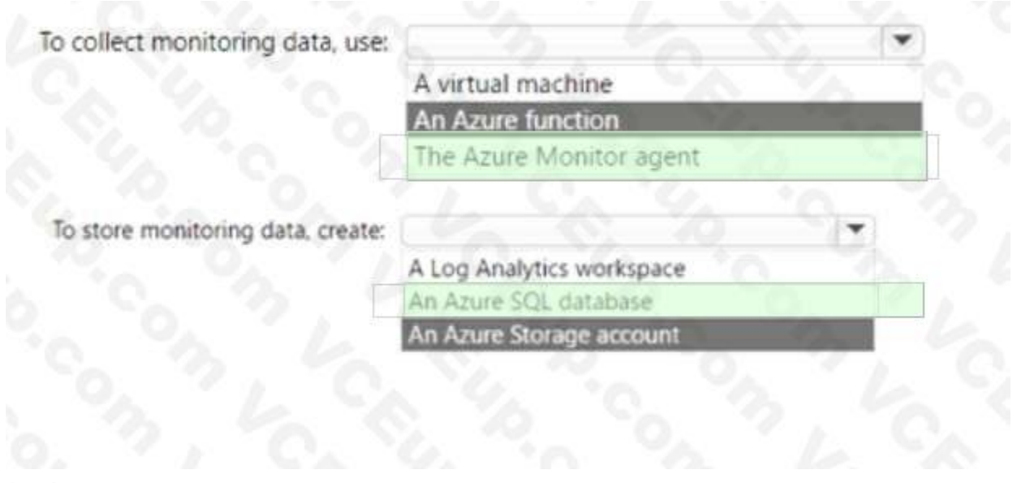

Section:
Explanation:
https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paasoverview?view=azuresql

QUESTION 27
You have an Azure subscription that contains an Azure SQL database named SQL1.
SQL1 is in an Azure region that does not support availability zones.
You need to ensure that you have a secondary replica of SQLI in the same region.
What should you use?

A. log shipping
B. auto-failover groups
C. active geo-replication
D. Microsoft SQL Server failover clusters

Correct Answer: C
Section:

QUESTION 28
Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.
You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.
What should you use?

A. Azure Migrate
B. Data Migration Assistant (DMA)
C. Azure Data BOX
D. Azure Database Migration Service

Correct Answer: C
Section:
Explanation:
https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-withMicrosoft-Azure-Migrate

QUESTION 29
DRAG DROP
You create a new Azure SQL managed instance named SQL1 and enable Database Mail extended stored.
You need to ensure that SQ Server Agent jobs running on SQL 1 can notify when a failure Occurs.
Which three actions should you perform in sequence 7 TO answer. move the appropriate actions from the list Of actions to answer area and arrange them in correct order.

Select and Place:

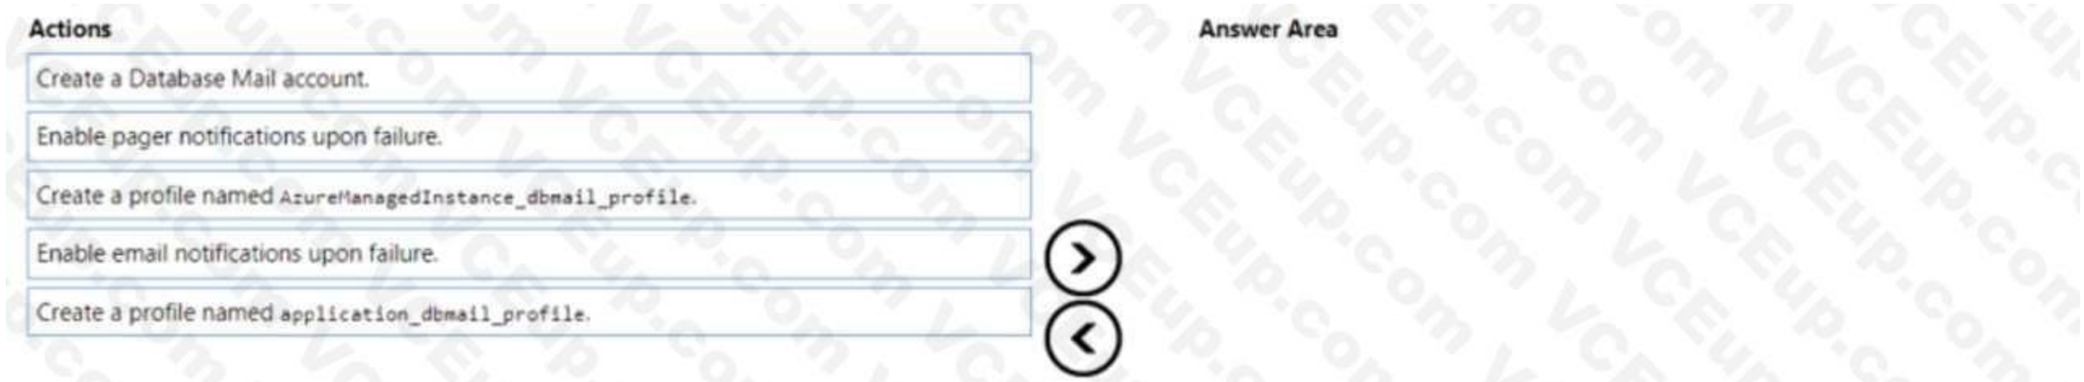

Correct Answer:

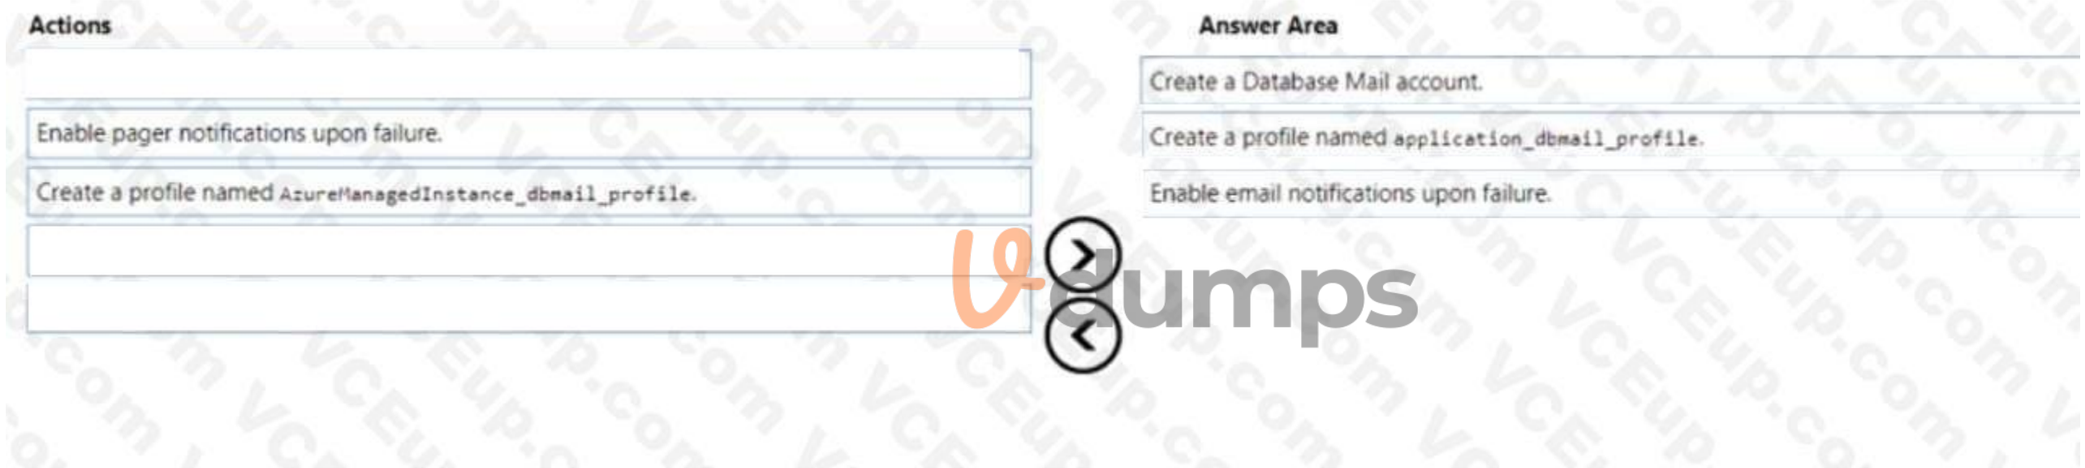

Section:
Explanation:

QUESTION 30
DRAG DROP
You have an Azure subscription that contains an Azure SQL database named SQLDb1. SQLDb1 contains a table named Table1. You plan to deploy an Azure web app named webapp1 that will export rows in Table1 that have changed.
You need to ensure that webapp1 can identity the changes to Table'. The solution must meet the following requirements:
• Minimize compute times.
• Minimize storage.
Which three actions should you perform in sequence? To answer. move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

Actions

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

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON
```

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

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_table
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_db
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING
```

Answer Area

Correct Answer:

Actions

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

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_table
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
EXEC sys.sp_cdc_enable_db
```

Answer Area

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER DATABASE SQLDb1 SET CHANGE_TRACKING = ON
```

Connect to SQLDb1 and run the following Transact-SQL statement.

```
ALTER TABLE dbo.Table1 ENABLE CHANGE_TRACKING
```

From webapp1, connect to SQLDb1, obtain the initial dataset, and r
function.

Section:
Explanation:

QUESTION 31
You have an Azure SQL database named DBI that contains a nonclustered index named index1.
End users report slow queries when they use index1.
You need to identify the operations that are being performed on the index.
Which dynamic management view should you use?

A. `sys.dm_exec_query_plan_stats`

B. `Sys.dm_db_index_physical_stats`

C. `Sys.dm_db_index_operational_stats`

D. `Sys.dm_db_index_useage_stats`

A. Option A
B. Option B
C. Option C

D. Option D

Correct Answer: D
Section:

QUESTION 32
You have a Microsoft SQL Server 2017 server.
You need to migrate the server to Azure. The solution must meet the following requirements:
• Ensure that the latest version of SQL Server is used.
• Support the SQL Server Agent service.
Minimize administrative effort.
What should you use?

A. SQL Server on Azure Virtual Machines
B. Azure SQL Database
C. an Azure SQL Database elastic pool
D. Azure SQL Managed Instance

Correct Answer: A
Section:

QUESTION 33
You have two on-premises Microsoft SQL Server 2019 instances named SQL1 and SQL2.
You need to migrate the databases hosted on SQL 1 to Azure. The solution must meet the following requirements:
The service that hosts the migrated databases must be able to communicate with SQL2 by using linked server connections. Administrative effort must be minimized.
What should you use to host the databases?

A. a single Azure SQL database
B. an Azure SQL Database elastic pool
C. SQL Server on Azure Virtual Machines
D. Azure SQL Managed Instance

Correct Answer: D
Section:

QUESTION 34
HOTSPOT
You have an Azure SQL database named DB1 that contains a table named Orders. The Orders table contains a row for each sales order. Each sales order includes the name of the user who placed the order. You need to implement row-level security (RLS). The solution must ensure that the users can view only their respective sales orders. What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

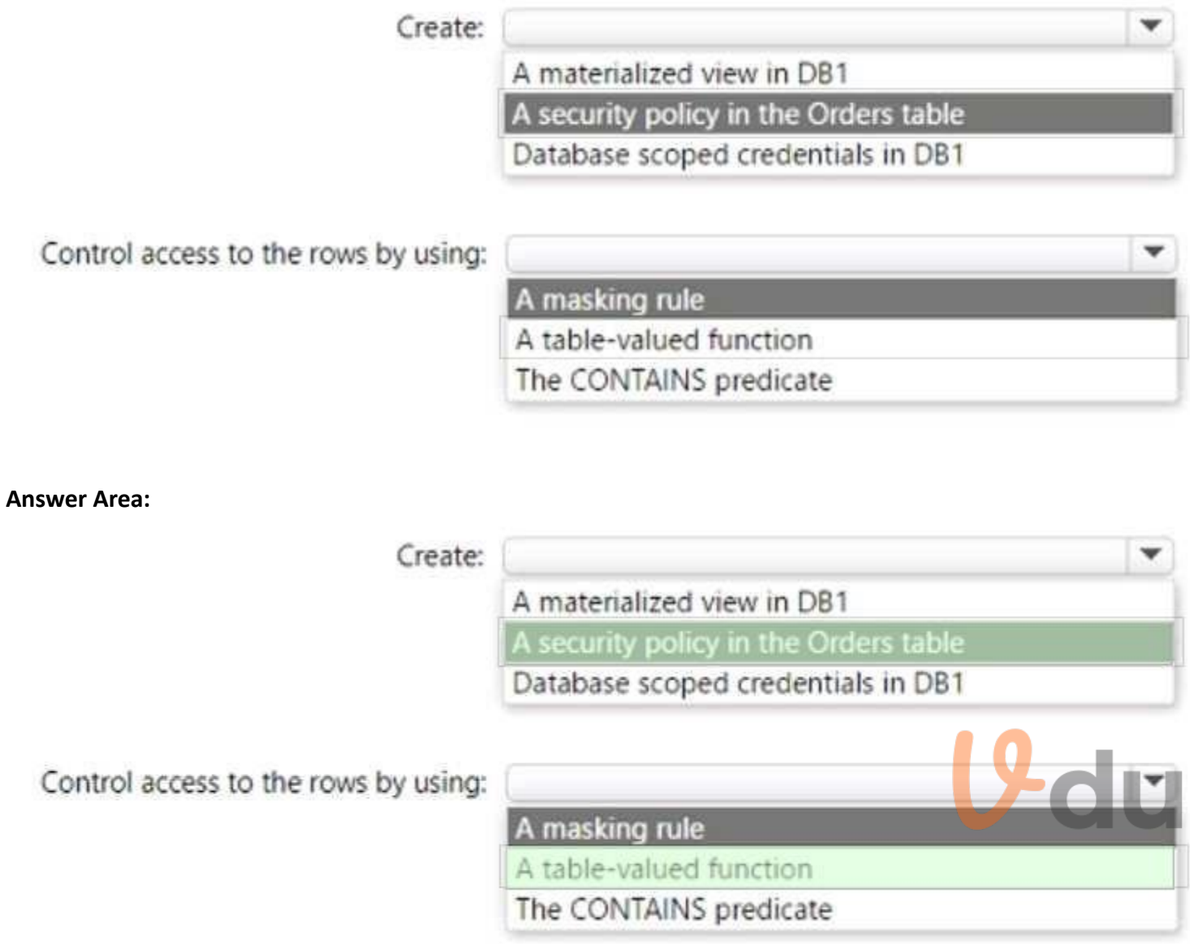

Answer Area:

Section:
Explanation:

QUESTION 35
HOTSPOT
You have a SQL Server on Azure Virtual Machines instance that hosts a database named Db1.
You need to configure the autogrow and autoshrink settings for DB1.
Which statements should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

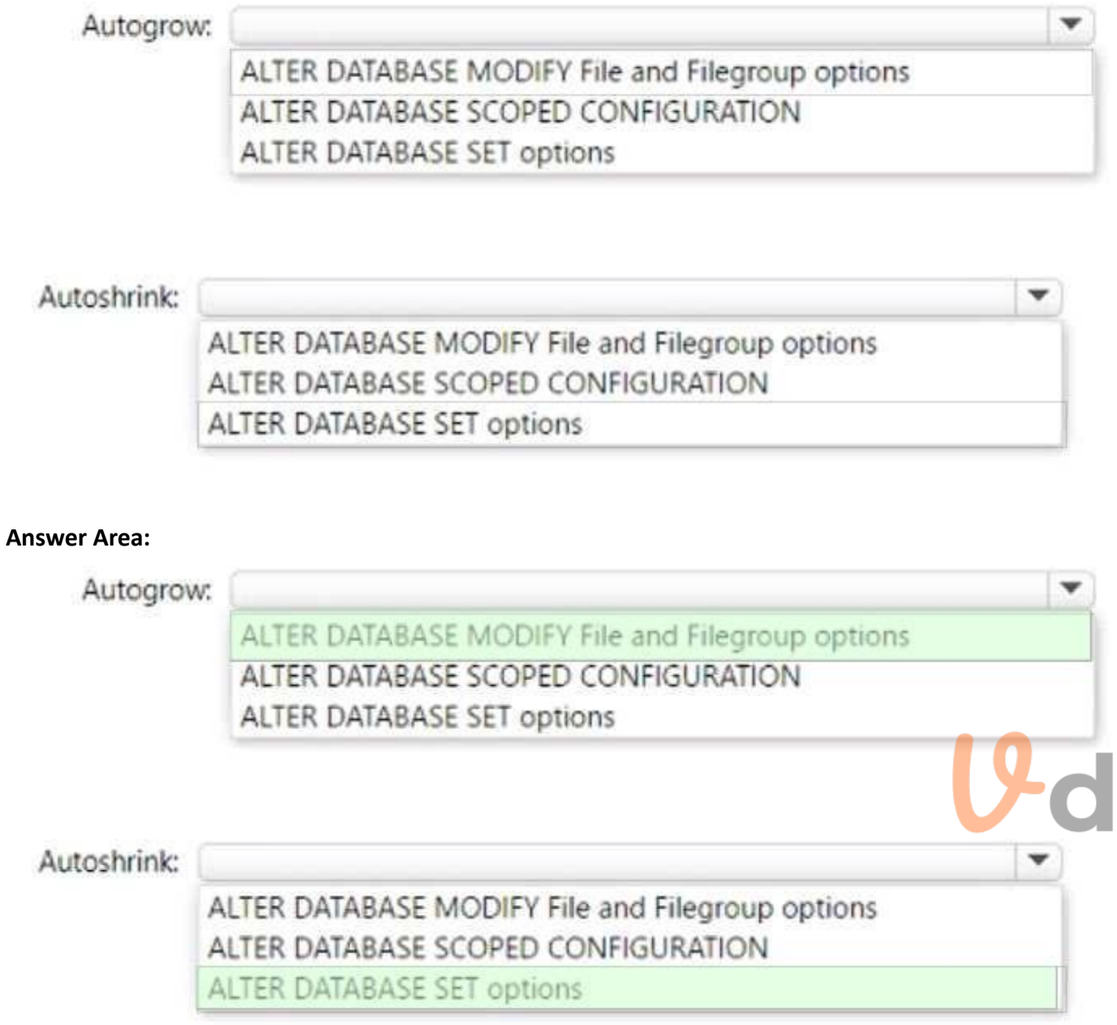

Answer Area:

Section:
Explanation:
https://learn.microsoft.com/en-us/troubleshoot/sql/admin/considerations-autogrow-autoshrink

QUESTION 36
HOTSPOT
You have an Azure SQL logical server.
You run the following script.

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

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

NOTE: Each correct selection is worth one point

Hot Area:

Statements	Yes	No
The `Orders` table will allow only rows to be inserted.	○	○
To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.	○	○
To ensure that a timestamp is added to each row in the `Orders` table, the `GENERATED ALWAYS`	○	○

Answer Area:

Statements	Yes	No
The `Orders` table will allow only rows to be inserted.	○	○
To create additional tables in the Sales database, the `LEDGER = ON` parameter must be used.	○	○
To ensure that a timestamp is added to each row in the `Orders` table, the `GENERATED ALWAYS`	○	○

Section:
Explanation:

QUESTION 37
HOTSPOT
You have an Azure subscription that is linked to an Azure AD tenant named contoso.com. The subscription contains an Azure SQL database named SQL 1 and an Azure web named app1. App1 has the managed identity feature enabled. You need to create a new database user for app1.

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

Hot Area:

Answer Area:

Section:
Explanation:
https://learn.microsoft.com/en-us/azure/app-service/tutorial-connect-msi-sqldatabase?tabs=windowsclient%2Cef%2Cdotnet

QUESTION 38
HOTSPOT
You need to use an Azure Resource Manager ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance. The solution must maximize disk I/O performance for the SQL Server database and log files How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

```
"variables": (
    "dataDisks": {
      "caching":  "dataDiskCount": 8,"logDisksCount": 1,
…
}
"resources": [
      …
osDisk": {
 …
,"copy": [
 {"name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
  "input": {"lun": "[copyIndex('dataDisks')]","createOption":"empty",
   "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),parameters('dataDiskCount')),
    variables('dataDisks').caching )]","diskSizeGB": 1023,
```

None
ReadOnly
ReadWrite

None
ReadOnly
ReadWrite

Answer Area:

Section:
Explanation:

QUESTION 39
You manage 100 Azure SQL managed instances located across 10 Azure regions.
You need to receive voice message notifications when a maintenance event affects any of the 10 regions. The solution must minimize administrative effort. What should you do?

A. From the Azure portal, create a service health alert.
B. From the Azure portal, create an Azure Advisor operational excellence alert.
C. From Microsoft SQL Server Management Studio (SSMS), configure a SQL Server agent job.
D. From the Azure portal, configure an activity log alert.

Correct Answer: C
Section:

QUESTION 40
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:

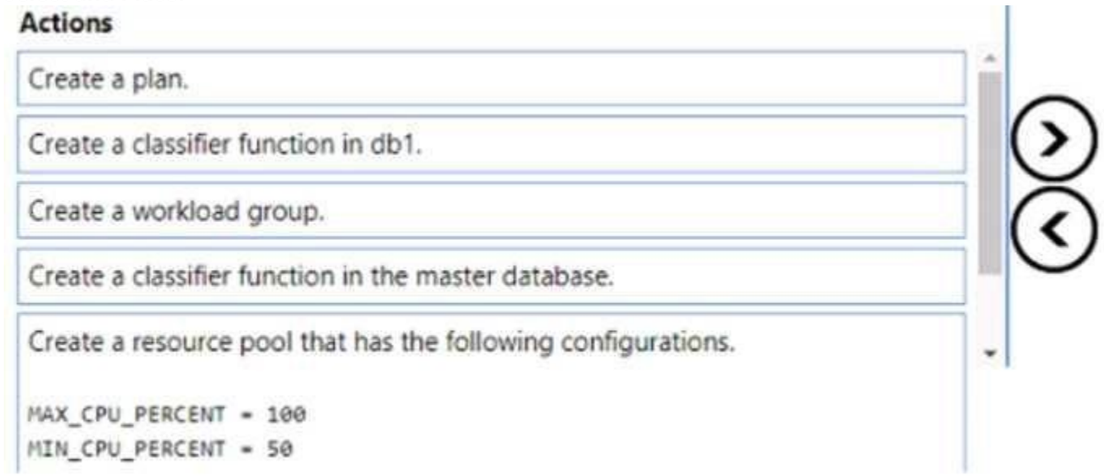

Correct Answer:

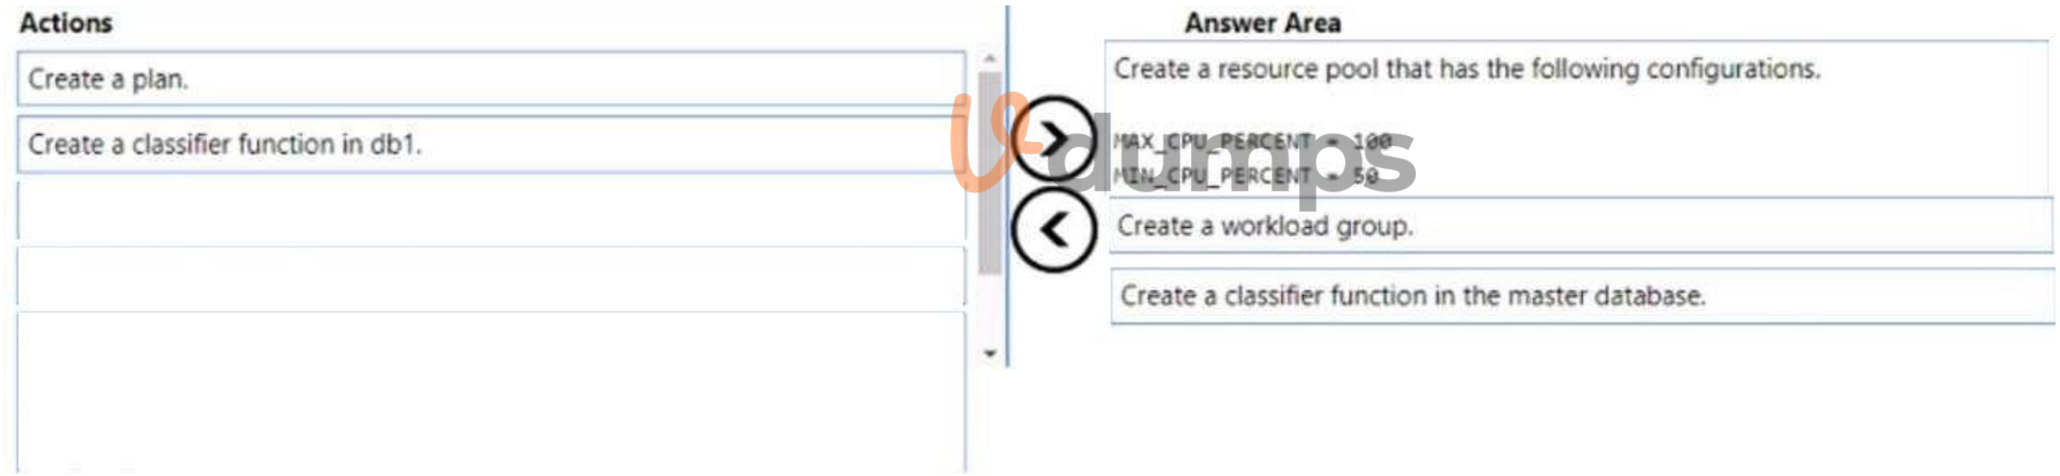

Section:
Explanation:

QUESTION 41
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:

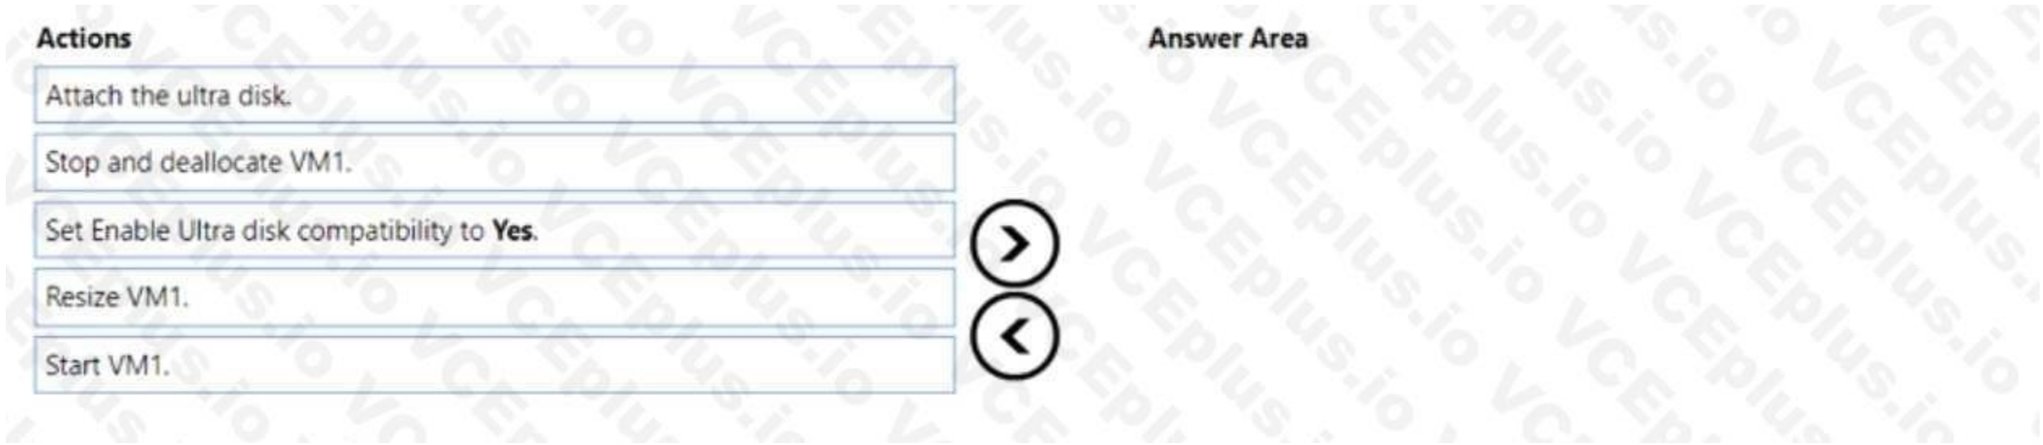

Correct Answer:

Section:
Explanation:

QUESTION 42
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:

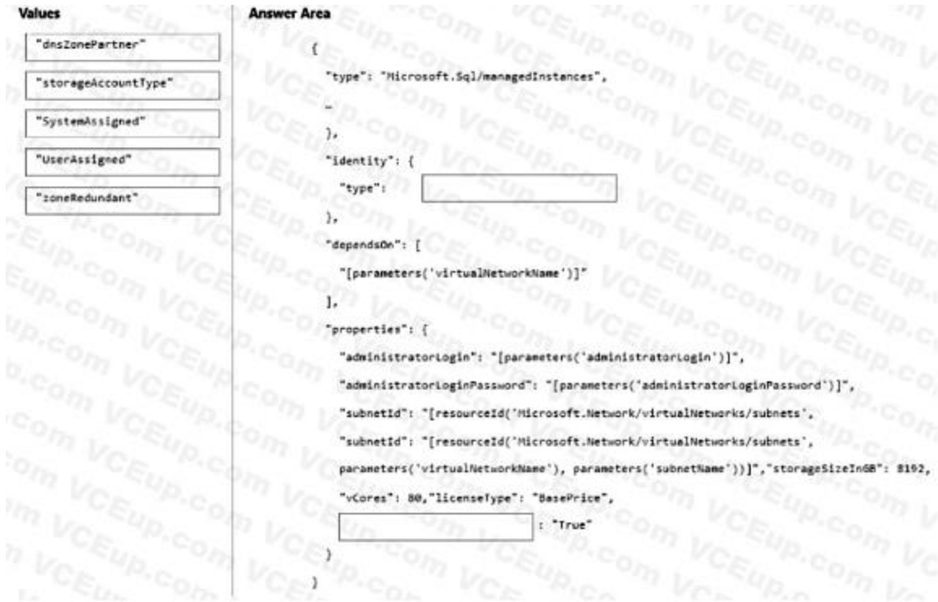

Correct Answer:

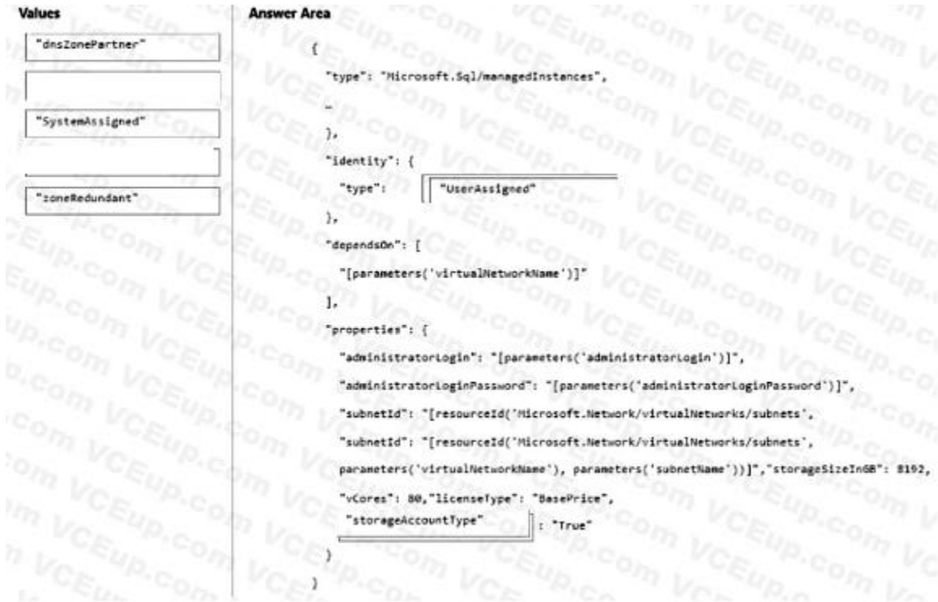

Section:
Explanation:

QUESTION 43
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

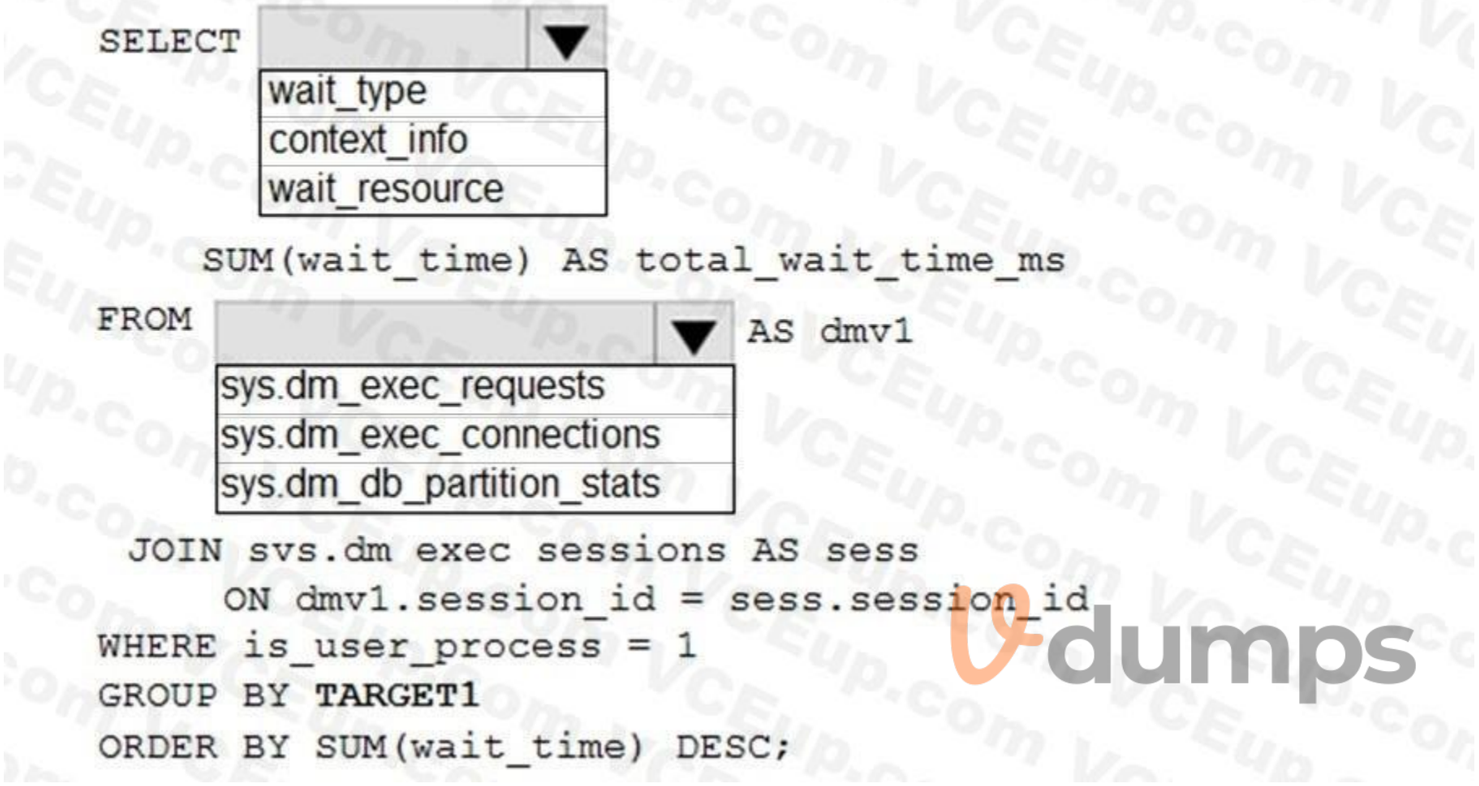

Answer Area:

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 44
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

```
SELECT SUM(size/ [16.0 | 128.0 | 512.0] - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB

FROM [sys.database_files | sys.resource_stats | sys.dm_db_resource_stats]

GROUP BY type_desc

HAVING type_desc = 'ROWS';
```

Answer Area:

Answer Area

```
SELECT SUM(size/ [16.0 | 128.0 | 512.0] - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0)
AS DatabaseDataSpaceAllocatedUnusedInMB

FROM [sys.database_files | sys.resource_stats | sys.dm_db_resource_stats]

GROUP BY type_desc

HAVING type_desc = 'ROWS';
```

Section:
Explanation:

QUESTION 45
HOTSPOT
You configure backup for an Azure SQL database as shown in the following exhibit.

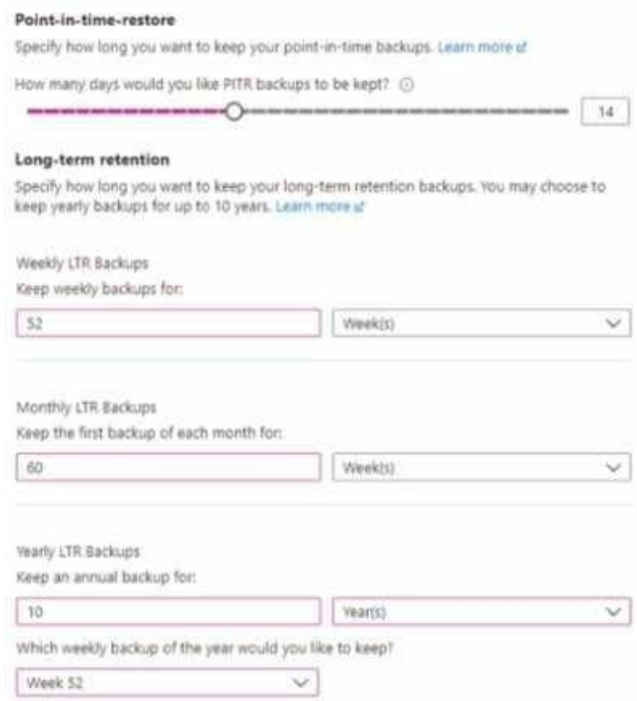

Use the drop-down menus to select the answer choice the completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

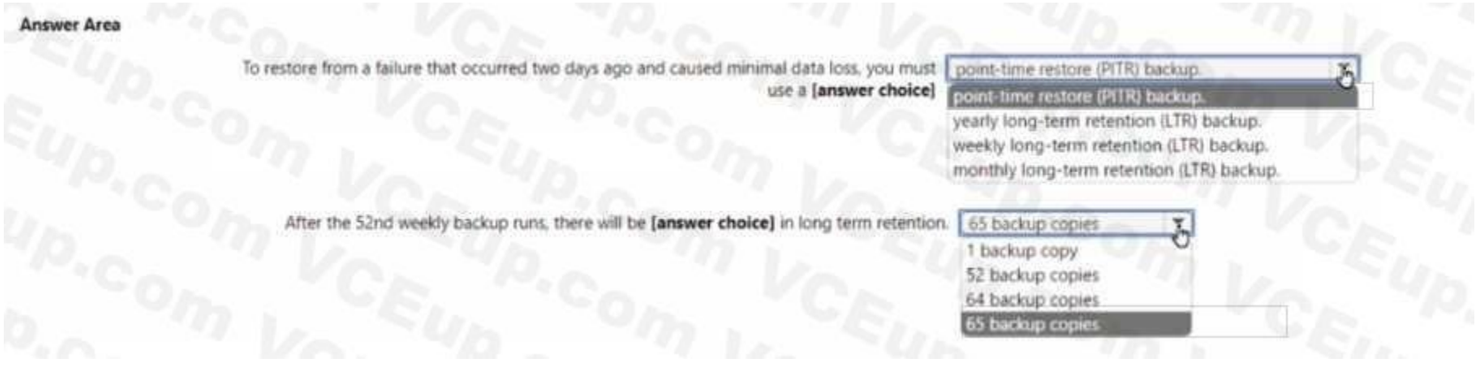

Answer Area:

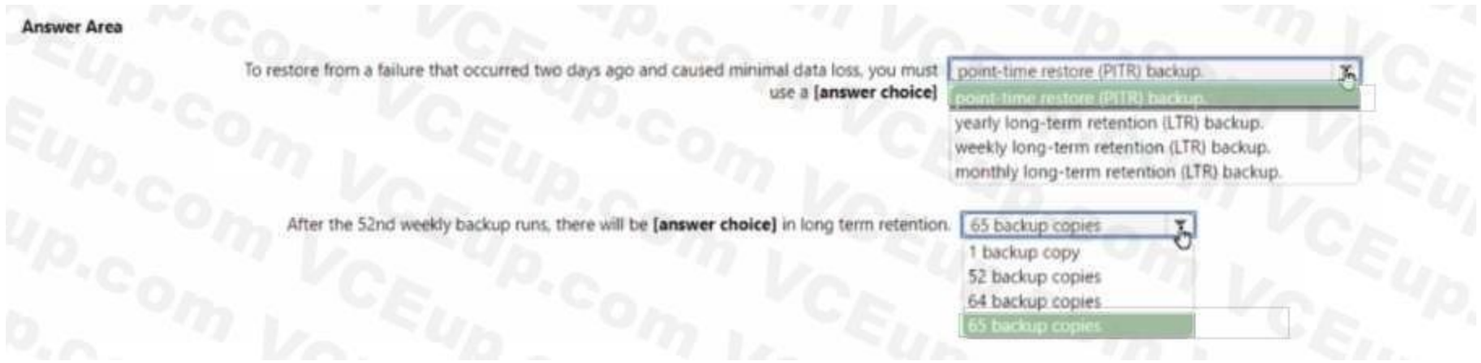

Section:
Explanation:

QUESTION 46
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 47
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 48
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 49
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:

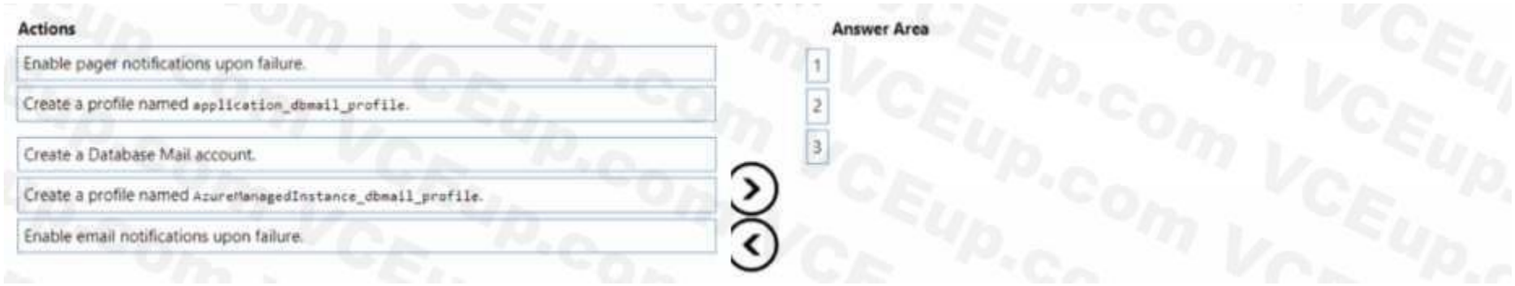

Correct Answer:

Section:
Explanation:

QUESTION 50
You have An Azure SQL managed instance.
You need to configure the SQL Server Agent service to email job notifications.
Which statement should you execute?

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

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

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

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

Correct Answer: B
Section:

QUESTION 51
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 52
Task 1
In an Azure SQL database named db1, you need to enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.SalesOrderHeader table.

A. See the explanation part for the complete Solution

Correct Answer: A
Section:
Explanation:
To enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.SalesOrderHeader table in db1, you can use the following Transact-SQL script:
-- Connect to the Azure SQL database named db1
USE db1;
GO
-- Enable page compression on the clustered index
ALTER INDEX PK_SalesOrderHeader_SalesOrderID ON SalesLT.SalesOrderHeader
REBUILD WITH (DATA_COMPRESSION = PAGE);
GO
This script will rebuild the clustered index with page compression, which can reduce the storage space and improve the query performance
The script solution consists of three parts:
The first part 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_SalesOrderID ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO. This part enables page compression on the clustered index named PK_SalesOrderHeader_SalesOrderID, which is defined on the SalesLT.SalesOrderHeader table. 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 53
HOTSPOT
You have an Azure SQL database named DB1.
You have 10 Azure virtual machines that connect to a virtual network subnet named Subnet 1.
You need to implement a database-level firewall that meets the following requirements:
* Ensures that only the 10 virtual machines can access DB1
* Follows the principle of least privilege
How should you configure the firewall rule, and how should you establish network connectivity from the virtual machines to DB1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Section:
Explanation: